



NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

**SIMULTANEOUS CONNECTION MANAGEMENT AND
PROTECTION IN A DISTRIBUTED MULTILEVEL SECURITY
ENVIRONMENT**

by

Joseph D. Sears

September 2004

Thesis Advisor:
Co-Advisor:

Cynthia E. Irvine
Thuy D. Nguyen

Approved for public release; distribution is unlimited

THIS PAGE INTENTIONALLY LEFT BLANK

| | | | | |
|--|---|--|--|--|
| REPORT DOCUMENTATION PAGE | | | <i>Form Approved OMB No. 0704-0188</i> | |
| Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503. | | | | |
| 1. AGENCY USE ONLY (Leave blank) | | 2. REPORT DATE September 2004 | 3. REPORT TYPE AND DATES COVERED Master's Thesis | |
| 4. TITLE AND SUBTITLE: Simultaneous Connection Management and Protection in a Multilevel Security Environment | | | 5. FUNDING NUMBERS | |
| 6. AUTHOR(S) Joseph D. Sears | | | | |
| 7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000 | | | 8. PERFORMING ORGANIZATION REPORT NUMBER | |
| 9. SPONSORING /MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A | | | 10. SPONSORING/MONITORING AGENCY REPORT NUMBER | |
| 11. SUPPLEMENTARY NOTES: The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government. | | | | |
| 12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited | | | 12b. DISTRIBUTION CODE | |
| 13. ABSTRACT (maximum 200 words) The Naval Postgraduate School Center for Information Systems Security Studies and Research (CISR) is designing and developing a distributed multilevel secure (MLS) network known as the Monterey Security Architecture (MYSEA). MYSEA will permit the delivery of unmodified commercial off the shelf productivity software applications and data from a large number of single-level network domains (e.g., NIPRNET, SIPRNET, JWICS) to a trusted distributed operating environment that enforces MLS policies. The analysis and development of a communications framework necessary to support connections between multiple MLS servers and a set of high assurance network appliances supporting simultaneous access to multiple single level networks and their concurrent connection management is required to fulfill the goal of MYSEA. To enable this functionality, modifications to the existing MYSEA server, the development of a new high assurance communications security device - the Trusted Channel Module (TCM), and the implementation of a <i>trusted channel</i> between the MYSEA server and the TCM is required. This document specifies a framework for incorporating the high level design of the TCM, several trusted daemons and databases, plus the incorporation of a <i>trusted channel</i> protocol into MYSEA to enable a distributed MLS environment. | | | | |
| 14. SUBJECT TERMS Multilevel Security, Trusted Channel, High Assurance, Distributed MLS Network | | | 15. NUMBER OF PAGES 146 | |
| | | | 16. PRICE CODE | |
| 17. SECURITY CLASSIFICATION OF REPORT Unclassified | 18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified | 19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified | 20. LIMITATION OF ABSTRACT UL | |

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited.

**SIMULTANEOUS CONNECTION MANAGEMENT AND PROTECTION IN A
DISTRIBUTED MULTILEVEL SECURITY ENVIRONMENT**

Joseph D. Sears
Lieutenant, United States Navy
B.A., University of Kentucky, 1989

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN COMPUTER SCIENCE

from the

**NAVAL POSTGRADUATE SCHOOL
September 2004**

Author: Joseph D. Sears

Approved by: Cynthia E. Irvine, Ph.D.
Thesis Advisor

Thuy D. Nguyen
Co-Advisor

Peter J. Denning, Ph.D.
Chairman, Department of Computer Science

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

The Naval Postgraduate School Center for Information Systems Security Studies and Research (CISR) is designing and developing a distributed multilevel secure (MLS) network known as the Monterey Security Architecture (MYSEA). MYSEA will permit the delivery of unmodified commercial off the shelf productivity software applications and data from a large number of single-level network domains (e.g., NIPRNET, SIPRNET, JWICS) to a trusted distributed operating environment that enforces MLS policies. The analysis and development of a communications framework necessary to support connections between multiple MLS servers and a set of high assurance network appliances supporting simultaneous access to multiple single level networks and their concurrent connection management is required to fulfill the goal of MYSEA. To enable this functionality, modifications to the existing MYSEA server, the development of a new high assurance communications security device - the Trusted Channel Module (TCM), and the implementation of a *trusted channel* between the MYSEA server and the TCM is required. This document specifies a framework for incorporating the high level design of the TCM, several trusted daemons and databases, plus the incorporation of a *trusted channel* protocol into MYSEA to enable a distributed MLS environment.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

| | | |
|------------|---|-----------|
| I. | INTRODUCTION..... | 1 |
| A. | BACKGROUND | 1 |
| 1. | MYSEA Historical Overview | 1 |
| 2. | Requirement to Expand Single Level Network Capability | 4 |
| 3. | Distributed MLS Network | 5 |
| B. | COMMON CRITERIA IMPACT ON HIGH ASSURANCE DEVELOPMENT | 8 |
| 1. | PalME Project | 8 |
| 2. | Trusted Channel Module Development | 9 |
| 3. | Terminology | 9 |
| C. | CHAPTER OVERVIEW | 9 |
| 1. | Introduction | 9 |
| 2. | MYSEA Connected Single Level Network Management Framework | 9 |
| 3. | Protected Communications Channel Protocol | 9 |
| 4. | The MYSEA Trusted Channel Module | 10 |
| 5. | Conclusion and Future Work | 10 |
| D. | APPENDIX OVERVIEW | 10 |
| 1. | Appendix A: Connected Single Level Network Systems Requirements Document | 10 |
| 2. | Appendix B: Protected Communications Channel Protocol Requirements Document | 10 |
| 3. | Appendix C: Trusted Channel Management Requirements Document | 10 |
| II. | MYSEA CONNECTED SINGLE LEVEL NETWORK MANAGEMENT FRAMEWORK..... | 11 |
| A. | CONNECTED SINGLE LEVEL NETWORK ARCHITECTURE COMPONENTS..... | 12 |
| 1. | MYSEA Server | 13 |
| 2. | Trusted Channel Module | 15 |
| 3. | Networking Device | 15 |
| 4. | Type I Encryption Device | 16 |
| B. | CONNECTED SINGLE LEVEL NETWORK SOFTWARE MODULES | 16 |
| 1. | MYSEA Server | 17 |
| a. | Trusted Channel Server | 18 |
| b. | Secure Connection Server | 18 |
| c. | Trusted Channel Database | 18 |
| d. | Secure Connection Inbound Database | 19 |
| e. | Secure Connection Outbound Database | 19 |
| f. | Protected Communications Channel Protocol handler | 19 |
| 2. | Trusted Channel Module | 20 |

| | | | |
|-------------------|-----------|---|------------|
| | <i>a.</i> | <i>Trusted Channel Server</i> | <i>21</i> |
| | <i>b.</i> | <i>Network Address Translation Server</i> | <i>21</i> |
| | <i>c.</i> | <i>Protected Communications Channel Protocol.....</i> | <i>21</i> |
| III. | | PROTECTED COMMUNICATIONS CHANNEL PROTOCOL..... | 23 |
| | A. | OVERVIEW | 23 |
| | B. | IPSEC PROTOCOL..... | 24 |
| | | 1. Security Policy Database | 24 |
| | | 2. Security Association Database | 25 |
| | | 3. Modes of Operation | 25 |
| | | 4. IPsec Security Protocols | 26 |
| | | 5. Security Association and Key Management | 27 |
| | | 6. IPsec Cryptographic Algorithms..... | 28 |
| | | 7. IPsec Placement..... | 28 |
| | C. | RESIDUAL RISK | 29 |
| | D. | SUMMARY | 30 |
| IV. | | THE MYSEA TRUSTED CHANNEL MODULE..... | 31 |
| | A. | TOE SECURITY ENVIRONMENT..... | 31 |
| | | 1. Assumptions | 31 |
| | | 2. Threats | 32 |
| | | 3. Organizational Security Policies..... | 33 |
| | B. | SECURITY OBJECTIVES | 35 |
| | | 1. Security Objectives for the TOE | 35 |
| | | 2. Security Objectives for the Environment | 38 |
| V. | | FUTURE WORK AND CONCLUSION | 39 |
| | A. | FUTURE WORK..... | 39 |
| | | 1. Formal TCM Security Target..... | 39 |
| | | 2. TCM Fail-Over..... | 39 |
| | | 3. Single-Level Network User Identification and Authentication | 39 |
| | | 4. MYSEA Security Association Protocol..... | 39 |
| | | 5. MLS IP Encryptors..... | 40 |
| | | 6. MLS Server-to-MLS Server Connectivity..... | 41 |
| | | 7. Protected Communications Channel Residual Risk Analysis | 41 |
| | B. | CONCLUSION | 41 |
| APPENDIX A | | CONNECTED SINGLE LEVEL NETWORK SYSTEM REQUIREMENTS DOCUMENT..... | 45 |
| APPENDIX B | | PROTECTED COMMUNICATIONS CHANNEL PROTOCOL REQUIREMENTS DOCUMENT | 67 |
| APPENDIX C | | CONNECTED SINGLE LEVEL NETWORK TRUSTED CHANNEL MANAGEMENT REQUIREMENTS DOCUMENT..... | 87 |
| | | LIST OF REFERENCES | 119 |
| | | INITIAL DISTRIBUTION LIST | 123 |

LIST OF FIGURES

| | | |
|-----------|---|----|
| Figure 1. | Current MYSEA Architecture | 2 |
| Figure 2. | MYSEA Architecture Incorporating TCM | 6 |
| Figure 3. | MYSEA CSLN Architecture | 12 |
| Figure 4. | CSLN with Remote TCM Augmented with Type I Encryption | 13 |
| Figure 5. | STOP / Intel Pentium Architecture Relationship [9] | 14 |
| Figure 6. | Transport Mode IPsec Packet using ESP Protocol | 26 |
| Figure 7. | Tunnel Mode ESP IPsec Packet..... | 26 |
| Figure 8. | AH/ESP Tunnel Mode Protected Packet | 27 |

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

| | | |
|----------|--|----|
| Table 1. | Security Usage Assumptions | 32 |
| Table 2. | Anticipated Threats Against TOE..... | 33 |
| Table 3. | Organizational Security Policies Applicable to the TOE | 34 |
| Table 4. | TOE Security Objectives | 37 |
| Table 5. | Security Objectives for the Environment..... | 38 |

THIS PAGE INTENTIONALLY LEFT BLANK

ABBREVIATIONS, ACRONYMS AND DEFINITIONS

Abbreviations, Acronyms

| | |
|----------|---|
| BITW | Bump-in-the-Wire |
| BLP | Bell and LaPadula Security Model |
| COIN | Coalition Interoperability Network |
| COMPUSEC | Computer Security |
| COMSEC | Communications Security |
| CSLN | Connected Single-Level Network |
| DAC | Discretionary Access Control |
| DoD | Department of Defense |
| EAL | Evaluated Assurance Level |
| IA | Information Assurance |
| IKE2 | Internet Key Exchange Protocol, Version 2 |
| IP | Internet Protocol |
| JWICS | Joint Worldwide Intelligence Communications System |
| MAC | Mandatory Access Control |
| MLS | Multilevel Security |
| MYSEA | Monterey Security Architecture |
| NIPRNET | NonSecure Internet Protocol Router Network |
| PCC | Protected Communications Channel |
| PP | Protection Profile |
| RFC | Request for Comment |
| SA | Security Association |
| SCS | Secure Connection Server |
| SCIDB | Secure Connection Inbound Database |
| SCODB | Secure Connection Outbound Database |
| SIPRNET | SECRET Internet Protocol Router Network |
| SPD | Security Policy Database |

| | |
|-------|---|
| ST | Security Target |
| STOP | Secure Trusted Operating System |
| TCB | Trusted Computing Base |
| TCM | Trusted Channel Module |
| TCDB | Trusted Channel Database |
| TCS | Trusted Channel Server |
| TCSEC | Trusted Security Evaluation Criteria |
| TCX | Trusted Computing Exemplar |
| TNI | Trusted Network Interpretation of the TCSEC |
| TOE | Target of Evaluation |
| TPE | Trusted Path Extension |
| TSF | TOE Security Functions |
| TSP | TOE Security Policy |

Definitions

1. Connected Single Level Network (CSLN): The segment of the MYSEA architecture from the MYSEA server to its associated TCMs responsible for multiplexing a large number of single level networks into one MLS network interface on the MYSEA server.
2. Information Assurance: “Measures that protect and defend information and information systems by ensuring their availability, integrity, authentication, confidentiality, and non-repudiation. This includes providing for restoration of information systems by incorporating protection, detection, and reaction capabilities“ [1].
3. Inter-TSF Transfers: “Communicating data between the TOE and the security functions of other trusted IT products“ [2].
4. Inter-TSF Trusted Channel: “Provides for the secure communication of user or TSF data between the TOE and another trusted IT product“ [3].
5. Internal TOE Transfer: “Communicating data between separated parts of the TOE Security Function (SF): A part or parts of the TOE that have to be relied upon for enforcing a closely related subset of the rules from the TSP“ [2].

6. Reference Monitor: “The concept of an abstract machine that enforces TOE access control policies“ [2].
7. Sensitivity Level: The combined classification of data based upon its security or classification level and integrity level.
8. Target of Evaluation (TOE): “An IT product or system and its associated guidance documentation that is the subject of an evaluation“ [2].
9. TOE Security Functions (TSF): “A set consisting of all hardware, software, and firmware of the TOE that must be relied upon for the correct enforcement of the TSP“ [2].
10. TOE Security Functions Interface (TSFI): “A set of interfaces, whether interactive (man-machine interface) or programmatic (application programming interface), through which TOE resources are accessed, mediated by the TSF, or information is obtained from the TSF“ [2].
11. TOE Security Policy (TSP): “A set of rules that regulate how assets are managed, protected and distributed within a TOE“ [2].
12. Trusted channel: A means by which a TSF and a remote trusted IT product can communicate with necessary confidence to support the TSP [2].
13. Trusted Channel Module: Security device required to enable “Inter-TSF Trusted Channels” between the MYSEA server and its authorized single level network.
14. Trusted Path: “A means by which a user and a TSF can communicate with necessary confidence to support the TSP“ [2].
15. Trusted Network Interpretation (TNI): “Provides interpretations of the Trusted Computer Security Evaluation Criteria (TCSEC) appropriate for evaluating a network of computer and communication devices as a single system with a single Trusted Computing Base (TCB), called the Network Trusted Computing Base (NTCB), which is physically and logically partitioned among the components of the network“ [4].

THIS PAGE INTENTIONALLY LEFT BLANK

ACKNOWLEDGMENTS

I would like to thank my wife and children for their infinite patience while I pursued the completion of this thesis and my Masters in Computer Science. Without their personal support and sacrifice, I would not have been able to complete this endeavor.

I would also like to thank the faculty and staff of the Naval Postgraduate School. Although there are too many individuals to thank by name, the following individuals played a significant role in helping me to complete this thesis. Dr. Cynthia Irvine and Thuy Nguyen provided the guidance, insight and support, which made this thesis possible. Their breadth of knowledge and generous giving of time was invaluable to my successful completion of this work. Professors George Dinolt, Tim Levin, J.D. Fulp, and John Gibson all readily supported my research into policy, networking and IPsec issues. My special thanks go out to Dave Shifflett and Jean Khosalim, who spent numerous hours with me refining my understanding of the MYSEA architecture and its implementation on the DigitalNet, XTS-400. I would also like to express my appreciation to Capt. Francis Afinidad, USAF and Doctorial Candidate, who helped hammer out a number of networking and Department of Defense policy issues.

Finally, I would like to acknowledge the Fellowship granted to me by the Space and Naval Warfare Systems Command (SPAWAR), Code 2721. This work was performed under Funding Document Number N6600104WR00096, Job Order Number RCS9A. The financial support of SPAWAR facilitated the purchase of all equipment used to develop this thesis. More importantly, the Fellowship made possible the thesis research travel that was instrumental to my understanding several complex topics. My thanks go out to Lew Gutman, SSC San Diego, and his team for their comments and suggestions.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

This thesis documents the analysis and development of a communications framework between multiple multilevel security (MLS) servers and a set of high assurance network appliances to support simultaneous access to multiple single level networks and their concurrent connection management. This thesis describes the necessary modifications to an existing MLS architecture, the requirements for development of a new high assurance communications security device, and further analysis and validation of an existing secure communications protocol to enable the functionality described above. Furthermore, this thesis will expand on efforts to create a trusted distributed operating environment that enforces MLS policies while continuing to support unmodified commercial off the shelf productivity software. Ultimately, this communications framework will permit the delivery of applications and data from a large number of single-level network domains (e.g., NIPRNET, SIPRNET, JWICS) to a MLS-enabled enclave.

A. BACKGROUND

1. MYSEA Historical Overview

The Monterey Security Architecture (MYSEA) focuses MLS research at the Naval Postgraduate School Center for INFOSEC Studies and Research (CISR). These efforts began in the late 1990s in response to unfulfilled Department of Defense (DoD) MLS requirements and are captured in several documents [5] [6] [7] [8]. These unfulfilled MLS requirements have forced the DoD to support multiple system high networks, resulting in duplication of equipment, wasted manpower resources, exploding costs and reduced situational/tactical awareness. Research into MLS solutions may eventually enable MYSEA to provide the DoD with a deployable heterogeneous network solution, running at multiple classification levels while continuing to support commercial-off-the-shelf equipment and applications at the client end.

The current MYSEA is designed to provide a high assurance distributed MLS networking environment based upon a small set of high assurance security devices that are nearly transparent to the end user. This virtually transparent security architecture

permits clients to continue using the commercial-off-the-shelf (e.g., Microsoft Windows, Linux, and Intel) and government off the shelf operating systems and applications to which the end user has already been trained and is accustomed to operating. The current architecture utilizes three primary components to deliver its MLS environment: thin clients, an MLS Server and a Trusted Path Extension Device (TPE) as illustrated in Figure 1.

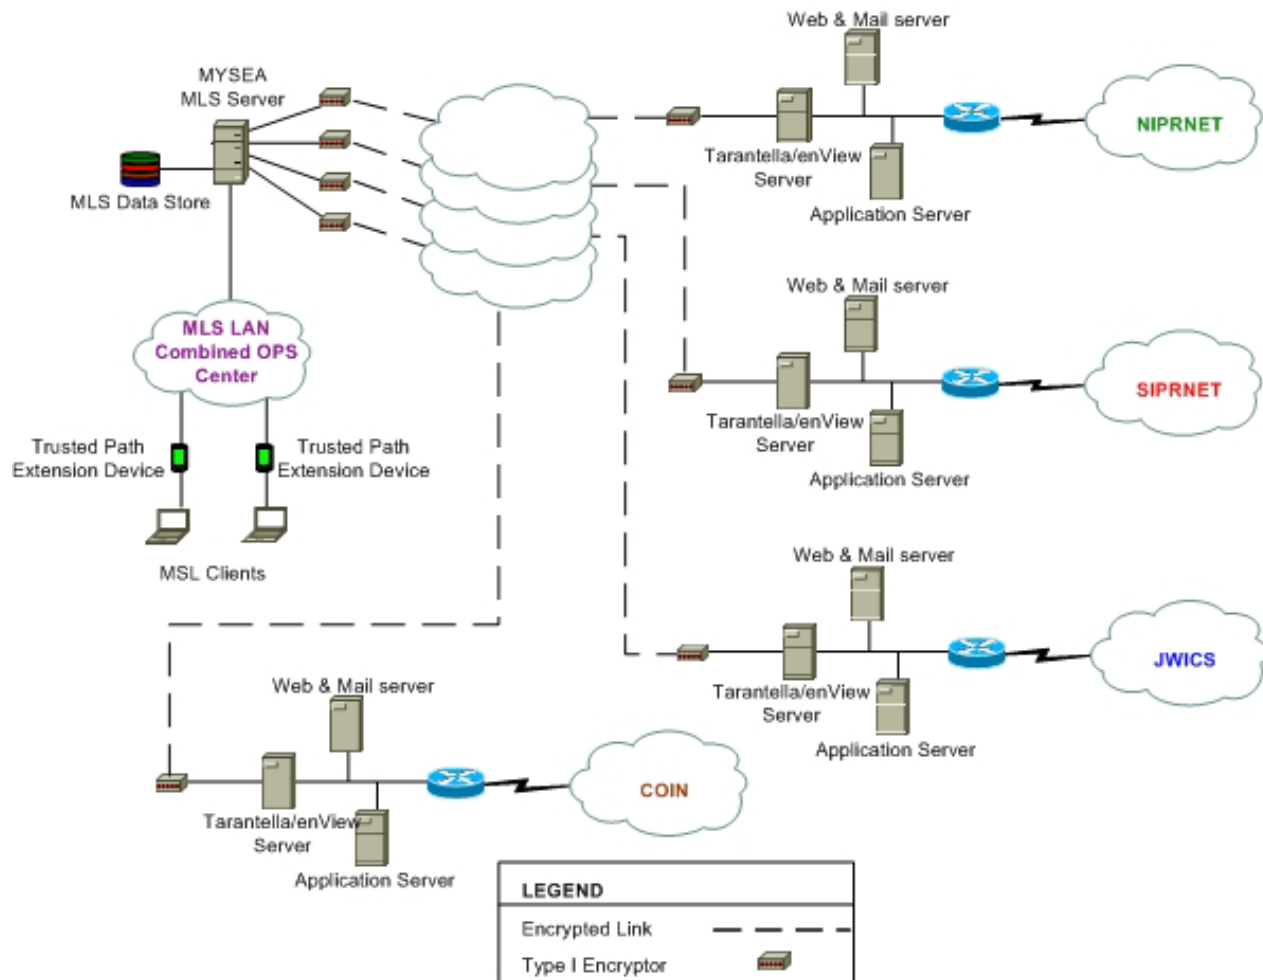


Figure 1. Current MYSEA Architecture

The prototype end user device is an Intel based platform running a Microsoft Windows XP or Linux diskless thin client, which is loaded with a standard commercial and government off the shelf package of applications at each user log-in. The client will have the capability to transition through multiple *sensitivity levels* of information with the

aid of the TPE. The client will remain in a diskless configuration so that no long-term memory storage is resident on the client after each transition.

The MYSEA server is a DigitalNet XTS-400 MLS server running the Secure Trusted Operating Program (STOP) 6.1 [9], previously evaluated at class B3 under the Trusted Computer System Evaluation Criteria (TCSEC) [10]. This operating system is currently undergoing a Common Criteria Evaluated Assurance Level (EAL) 5+ evaluation [2]. The STOP enforces MLS policies by using mandatory access controls (MAC) and discretionary policies with discretionary access controls (DAC). MAC and DAC policy enforcement is governed by the rules articulated in the Bell and LaPadula (BLP) security model [11] and in the Biba integrity model [12]. The BLP model shows how a secure state can be maintained by preventing unauthorized disclosure of information while the Biba integrity model shows how a secure state can be maintained by preventing the unauthorized modification of data.

Additionally, the STOP enforces DAC. This feature enables an owner of an object to explicitly control access to that object from other subjects. The STOP provides granular control over seven DAC modes including read and write. It also maintains permissions for the owner and groups authorized by the owner, as well as global permissions.

Responsible for interfacing each untrusted client to the MLS server and facilitating a true distributed MLS environment is the TPE device [5]. The TPE currently runs as a prototype on an open source operating system but will migrate to the Trusted Computing Exemplar (TCX) [13] high-security kernel under development at the Naval Postgraduate School. The TCX project aims to develop a high assurance separation kernel targeted for a Common Criteria EAL 7 evaluation. The TPE has a handheld form factor and extends the *trusted path* [3] of the MLS server through the MLS LAN. This functionality enables the user on an untrusted workstation to authenticate and negotiate a session at any authorized *sensitivity level* with the MYSEA server. Once logged-in at their authorized *sensitivity level*, a user will be able to access data at all *sensitivity levels* dominated by their current *sensitivity level*. In other words, if a user is logged-in at TOP

SECRET, they will have read and write access at TOP SECRET and read access to all *sensitivity levels* below TOP SECRET in compliance with the BLP and Biba models.

The MYSEA server has a limited number of physical network interfaces used to enhance its connectivity to MYSEA clients, LANs and networks. The MYSEA configured XTS-400 currently has four physical network interfaces to provide its connectivity with the possibility of expanding to sixteen network interfaces. One network interface is dedicated to the MLS LAN while the three remaining network interfaces are allocated to dedicated single level networks simulating current DoD architectures (e.g., JWICS, SIPRNET, NIPRNET, COIN). Each single-level network facilitates connectivity to an existing network infrastructure, thus permitting the use of these pre-existing networks to facilitate connectivity and interoperability with existing programs of record. Currently, each network interface is administratively assigned a dedicated *sensitivity level* corresponding to the requirements set forth for that single-level network (e.g., SIPRNET = SECRET/HIGH INTEGRITY). As a result, all communications traversing the network interface are assigned the proper *sensitivity level* by the STOP.

2. Requirement to Expand Single Level Network Capability

Although the possibility of connecting fifteen separate single-level networks to the MYSEA Server appears more than adequate, further research into the requirements for providing a true MLS environment reveals that number to be woefully inadequate. Williams and Day's critical analysis of "Sensitivity Labels and Security Profiles" [14] provides a mere glimpse into the nearly countless number of possible labeling combinations used by the DoD, its Coalition partners and contractors, and other U.S. Government entities. Likewise, Lipner's article on "Non-Discretionary Controls in Commercial Applications" [15] highlights the same issue in the private sector. With a multitude of existing single-level networks, the finite number of available MYSEA server network interfaces becomes a resource reserved only for the most significant single-level networks. This limitation impacts a true MLS environment by preventing the MLS hub from being fully connected to all levels of information. This inhibits the consolidation of information at one central location, negating the near real-time power of linking

repositories of critical information. The net result is a reduced ability to perform true fusion analysis and collaborative planning, which decreases situational and tactical awareness.

Recognition of this limitation generated a new requirement for MYSEA to expand the functionality of each network interface from a dedicated single-level network interface to a true MLS network interface capable of multiplexing multiple *sensitivity levels* of traffic. As the TPE extended the trusted path from the MYSEA Server to the MYSEA client, a new security device is also necessary to create a *trusted channel* [3] from the MYSEA server to each single-level network. With the creation of a *trusted channel*, it becomes possible to create an MLS network interface capable of managing and multiplexing a large number of single-level networks. This security device shall be known as the Trusted Channel Module (TCM), and its position in the MYSEA architecture is illustrated in Figure 2. Hereafter, the MYSEA segment that encompasses the *trusted channel* from the MYSEA server to the TCM will be known as a Connected Single Level Network (CSLN). The establishment of a *trusted channel* requires a Protected Communications Channel (PCC) Protocol to ensure that all communications between the MYSEA server and TCM are protected. Based upon the IPsec protocol suite, the PCC protocol will provide confidentiality, integrity and authenticity, satisfying the *trusted channel* requirements defined by the Common Criteria. Finally, implementation of the *trusted channel* will require additional processes running on the MYSEA server to manage the *trusted channel* communications.

3. Distributed MLS Network

Creating MLS network interfaces on the MYSEA server to provide multiple simultaneous *trusted channels* to multiple TCMs yields a new level of complexity for MYSEA. This new architecture produces an innovative distributed MLS network design. Although MYSEA pioneers new ground in its approach to creating a distributed MLS network, two other pioneering works exist to provide historical insight on developing secure distributed architectures. These two projects were developed in the late 1980s when the TCSEC [10] was the basis for assurance evaluation. Sufficient differences exist between TCSEC and Common Criteria terminology, particularly when trying to draw similarities between a Trusted Computing Base (TCB) and a Target of Evaluation (TOE)

Security Function (TSF). Additionally, the TCSEC Trusted Network Interpretation (TNI) incorporated the term Trusted Network Interface [4] for which there exists no Common Criteria equivalent. The TNI provided a guideline for evaluating a TCB distributed architecture of individually evaluated products as one complete system.

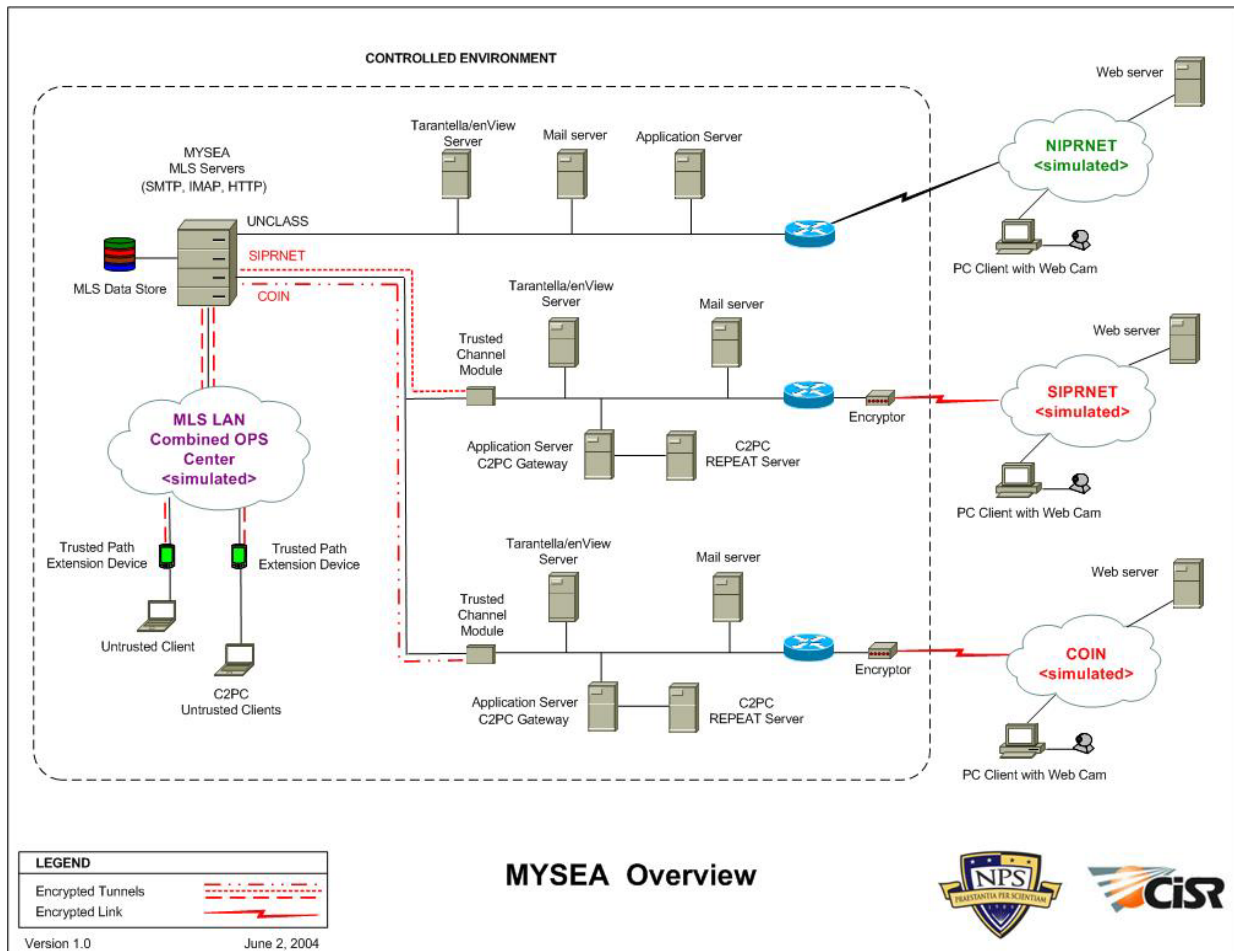


Figure 2. MYSEA Architecture Incorporating TCM

“The Architecture of a Distributed Trusted Computing Base,” [16] by Fellows et al., first introduced the underlying concepts concerning the successful implementation of a distributed MLS architecture. The authors explored five primary areas of focus when building a distributed MLS network. First, the system shall account for a “fragmented TCB (Trusted Computing Base) domain.” In other words, the distributed nature of all security devices within the MLS architecture complicates the ability to track and maintain *state* over all security functions within the system. Therefore, distributed systems shall

account for the delay and state transitions of the remote security components. Second, the architecture shall maintain “trusted paths between TCB components.” MYSEA maintains a trusted path between the MYSEA Server and its users through the TPE. The advent of the Common Criteria introduces new terminology to replace the term *trusted path* when describing data flows over distributed networks. The term *trusted channel* is introduced to differentiate the subtle difference between it and *trusted path*. For MYSEA, *trusted path* is interpreted to mean when the communication is from a human to a trusted system and *trusted channel* is interpreted to mean when the communication is between trusted systems [2]. Regardless of the terminology, the overarching principle is that an unforgeable link shall be maintained between the MLS server and other trusted devices within the network.

Third, the system shall invoke “trusted protocols” to ensure end-to-end security of the *trusted channel*. These “trusted protocols” typically use cryptography. Fourth, the system shall implement “hierarchical trusted computing bases.” This requirement maintains that at all times one central security device shall have a *reference monitor* capable of maintaining the integrity of the “multilevel secure environment.” Fifth, the system shall be constructed with “fault tolerance.” This requirement maintains that with any distributed MLS environment, if one of the security devices fails, it will “fail-secure” to prevent any compromise of the system or its data. Furthermore, fault tolerance may include the distribution of data and processing within the TCB to prevent the failure of one device from becoming a denial of service on the entire network

Fellows’ paper presents two further traits of distributed MLS architectures, known as “entelechy” and the “ Δ (change) property.” As defined by Fellows, “the entelechy component of the security policy states that a host may send or receive messages over a crypto connection only if it has current access to that crypto connection.” Entelechy is further explained as an important component of a distributed MLS architecture, because so much of the trusted security policy decision-making process is incorporated into trusted software. Therefore, entelechy becomes a critical factor in enforcing the overall security policy of the system. The “ Δ property” specifies that any security-related modification to the MLS distributed architecture shall only be made by trusted agents explicitly authorized to make said modifications.

Wiessman, in a paper entitled “BLACKER: Security for the DDN” [17], provided the second work which helped lay the foundations for a distributed MLS architecture. Wiessman first presented this paper in 1986, detailing the implementation of a secure host-to-host communication system for the Defense Data Network. Although the focus of the paper is based upon how best to build a TCSEC Class A1 level communications system, one notable trait instrumental for the success of a distributed MLS system is the capability to establish a *trusted path*. Weissman states that by blending both computer security (COMPUSEC) and communications security (COMSEC) principles, a “cryptographic seal” can be established between security devices, thus expanding the TCB across all security devices.

B. COMMON CRITERIA IMPACT ON HIGH ASSURANCE DEVELOPMENT

1. PalME Project

Development of high assurance systems requires meticulous planning, detailed analysis, and complete oversight of the entire system lifecycle process. The Common Criteria is the international standard for secure systems development. It provides a formalized set of functional and assurance security requirements necessary to achieve a desired level of evaluated security for any product. The Common Criteria permits tremendous flexibility for system designers to specify the level of security functionality and assurance their product requires and provides a detailed blueprint for how to achieve that security. Monika Vetterling and Guido Wimmel’s paper on “Secure Systems Development Based on the Common Criteria: The PalME Project” [18] details the benefits of secure systems development in conformance with the Common Criteria.

PalME demonstrated that incorporating a Common Criteria developmental framework from the conception of a project ensures analysis into the areas of threat, threat mitigation and residual risk. Additionally, the Common Criteria provides a set of accepted functional and assurance security requirements that, if implemented from the inception a project, will both focus and streamline product development by integrating security from the beginning versus backfilling security requirements as they arise.

2. Trusted Channel Module Development

Common Criteria development principles will be used in the development of the TCM. Currently, no specific *protection profile* (PP) exists to aid in this development. However, the “Consistency Instruction Manual for Development of U.S. Government Protection Profiles for use in Medium Robustness Environments” [19] establishes an initial baseline, which will be used for TCM development. Given this guideline, a high-level analysis of the threats, objectives and assumptions necessary to develop the TCM will be presented. Two additional documents will aid in the development of the TCM. The “ST for Cisco IOS/IPsec” [20] provides useful information for analyzing the fundamental network security elements necessary to establish a *trusted channel* based upon IPsec. Second, the “U.S. Government Protection Profile for Separation Kernels in Environments Requiring High Robustness – Version 0.621” [21] provides the separation kernel blueprint instrumental to the development of the TPE and TCM security devices.

3. Terminology

The following terms are used in accordance with the ISO definitions contained in ISO/IEC Directives Part 2, Rules for the structure and drafting of International Standards:” “Within normative text, the verbs ”shall”, ”should”, ”may”, and ”can” have the ISO standard meanings. [2]

C. CHAPTER OVERVIEW

1. Introduction

Chapter I defines the overall purpose of this thesis and provides the necessary background to understand distributed MLS architectures. It also provides an outline covering the remaining chapters and appendixes.

2. MYSEA Connected Single Level Network Management Framework

Chapter II defines the MYSEA CSLN management framework required to securely connect a large number of CSLNs to a single MLS network interface.

3. Protected Communications Channel Protocol

Chapter III reinforces the initial PCC findings denoted in “A Trusted Connection Framework for Multilevel Secure Local Area Networks” [5]. It will further define PCC requirements for establishing a *trusted channel*.

4. The MYSEA Trusted Channel Module

Chapter IV provides high level analysis of the TOE threat, objective and assumptions that a TCM must satisfy. It also discusses the development of the TCM in terms of Common Criteria guidelines.

5. Conclusion and Future Work

Chapter V summarizes the findings set forth in this thesis concerning distributed MLS architectures. Additionally, it provides a context for future work related to expanding the functionality of MYSEA.

D. APPENDIX OVERVIEW

1. Appendix A: Connected Single Level Network Systems Requirements Document

Appendix A critically analyzes all requirements necessary to incorporate CSLNs into MYSEA. These requirements will be broken down into three subsections covering the MYSEA server, TCM and CSLN operations applicable to both devices.

2. Appendix B: Protected Communications Channel Protocol Requirements Document

Appendix B provides a detailed analysis concerning the protocol requirements necessary to implement a *trusted channel* within the MYSEA architecture.

3. Appendix C: Trusted Channel Management Requirements Document

Appendix C provides a detailed analysis of the CSLN architecture required to support labeling of all CSLN connections. Detailed analysis will be provided on *trusted channel* management operations.

II. MYSEA CONNECTED SINGLE LEVEL NETWORK MANAGEMENT FRAMEWORK

The CSLN architecture encompasses all hardware, firmware and software necessary to create a *trusted channel* from the MYSEA server to the TCM. The *trusted channel* provides the foundation for all CSLN operations and is the principle mechanism for providing a distributed MLS architecture between the MYSEA server and the associated TCMs. The MYSEA server will act as a master in a master/slave relationship with each associated TCM. The MYSEA server will enforce all security policy decisions concerning the assigned *sensitivity level* of each CSLN and will maintain ultimate authority for permitting or denying all communications between itself and the CSLNs. This hierarchical CSLN architecture is an essential element for creating a distributed MLS environment according to Fellows [16]. It ensures that at all time, one central security device maintains a *reference monitor* capable of maintaining the integrity of the distributed MLS network environment.

The CSLN architecture will permit MYSEA to multiplex a large number of single-level networks, each operating at its own dedicated *sensitivity level* through one MLS network interface. The *trusted channel* will provide the MYSEA server with the underlying integrity mechanisms required to extract an implicit *sensitivity level* from all inbound CSLN connections to the MYSEA server and use that information to generate an explicit *sensitivity level* for all MAC based decisions. The *trusted channel* will also provide the MYSEA server with the necessary mechanisms to ensure all outbound connections destined to a CSLN are directed only to that CSLN.

As stated by both Fellows and Weissman, establishing a *trusted channel* is the principle element necessary to create a distributed MLS architecture. The *trusted channel* is responsible for permitting only authorized connections between the TCM and the MYSEA server. High level TCM requirements for establishing a *trusted channel* are presented in Chapter III.

This chapter will discuss how CSLN management functions are implemented in hardware, firmware and software. The high-level design requirements and specification data can be found in Appendix A [22] and C [23] of this thesis.

A. CONNECTED SINGLE LEVEL NETWORK ARCHITECTURE COMPONENTS

A CSLN consists of three components as depicted in Figure 5: 1) the MYSEA server; 2) the TCM; and 3) an untrusted networking device. The MYSEA server running on the XTS-400 is the central point of control mediating all accesses to data at different security levels. The TCM will serve as the primary interface between an authorized single-level network and the MYSEA server. A networking device will permit the multiplexing of a large number of TCMs into one MLS network interface on the MYSEA server.

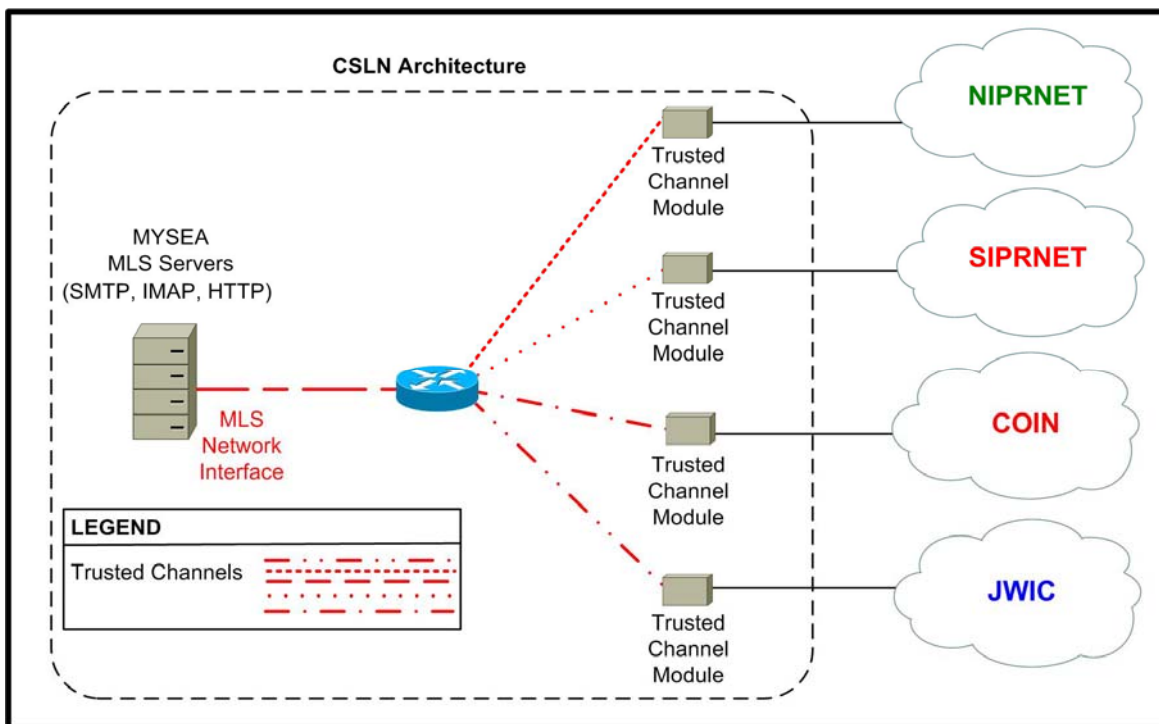


Figure 3. MYSEA CSLN Architecture

One additional set of components may be added to the CSLN architecture to provide additional COMSEC. Type I, NSA approved encryption devices could be used to cover CSLN communications links where the TCM is geographically separated from

the MYSEA server trusted physical environment. Although many configurations of a Type I modified architecture exist, one possible implementation is presented in Figure 4.

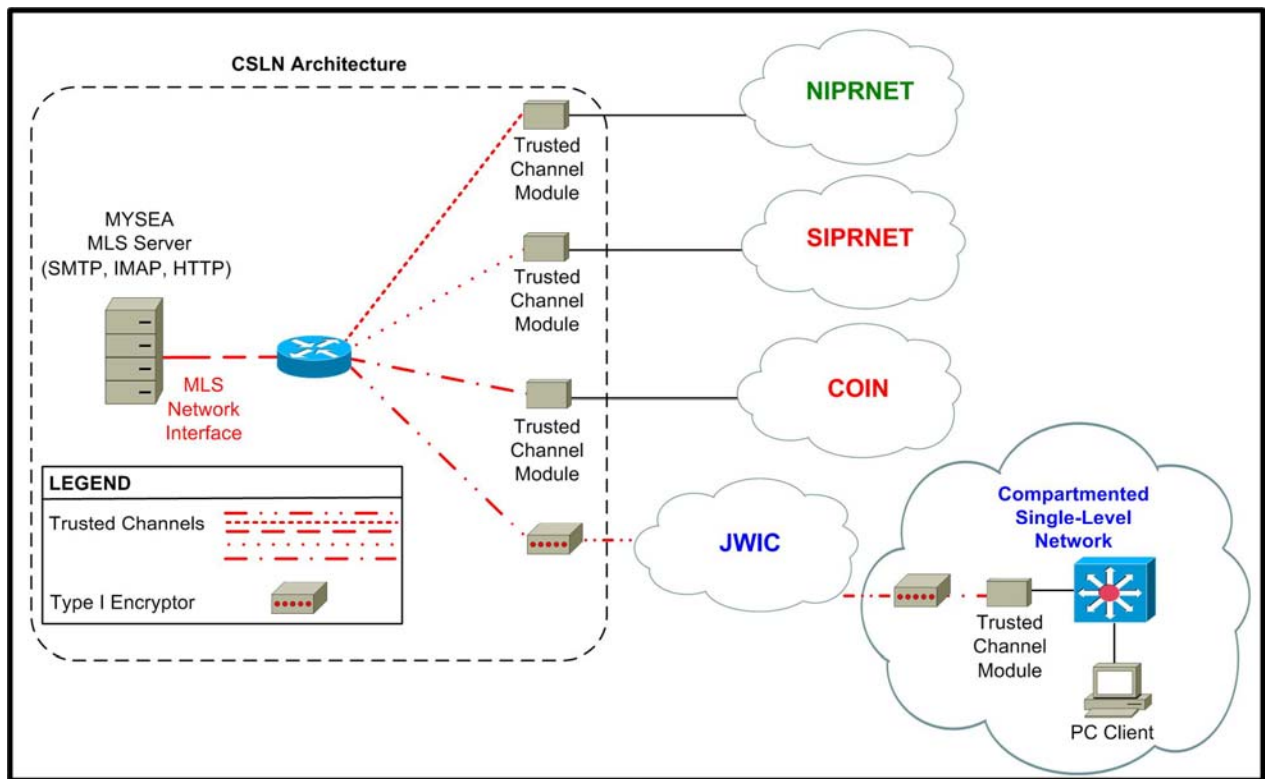


Figure 4. CSLN with Remote TCM Augmented with Type I Encryption

1. MYSEA Server

The DigitalNet XTS-400 hardware architecture and the STOP provide the security foundation for the MYSEA server. The XTS-400 is currently undergoing an EAL 5+ certification. This certification confers a high degree of confidence that the XTS-400 correctly enforces its TOE Security Policy (TSP) [2]. The XTS-400 leverages the *four domain architecture* of the Intel Pentium chip to fortify the STOP with hardware based security mechanisms [9]. Figure 5 illustrates the synergistic security relationship between the Pentium and STOP architectures. The x86 *four protection domain* architecture permits the separation of STOP trusted and untrusted operations at the hardware level.

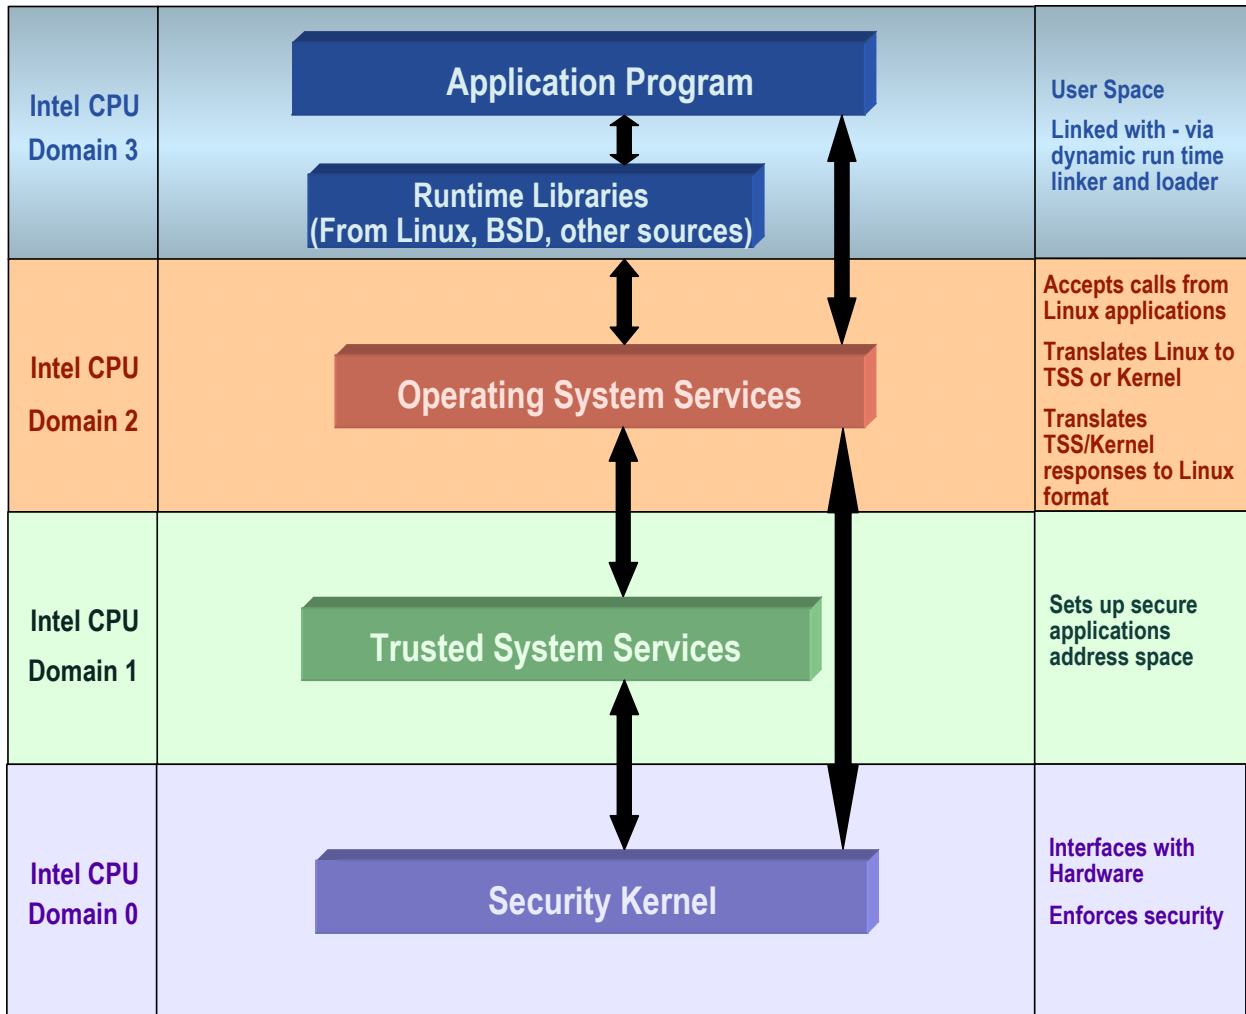


Figure 5. STOP / Intel Pentium Architecture Relationship [9]

Domain 0 is responsible for enforcing all security kernel operations and maintains the *reference monitor*. The *reference monitor* is “an abstract machine that enforces TOE access control policies [2] and is responsible for ensuring no software process with less privilege can run with kernel level privileges. All user level processes running in Domain 3 are forced to interact with the STOP through *Domain 2*. *Domain 2* provides an interface which mediates all user level access to the trusted domains of operation in the STOP. Combined with the STOP security functions, the Intel hardware domain separation between privilege levels provides a high degree of confidence that untrusted user level processes cannot escalate their privilege and thus tamper with STOP trusted processes for the purpose of subverting the system.

2. Trusted Channel Module

The TCM is the integral piece of hardware necessary to create a *trusted channel* through the MYSEA server MLS network interface to a single-level network. The TCM will be developed with the goal of achieving an EAL 6 certification. At present, it remains in the high level design phase. The TCM will also incorporate the Intel x86 four domain architecture. However, specific form factor hardware requirements remain unidentified and are left for *future work*.

3. Networking Device

The CSLN architecture will require an untrusted dumb hub or smart switching device to multiplex two or more CSLNs to a single MLS network interface on the MYSEA server. Although this device may yet have an EAL, the current generation of networking devices is not evaluated as high assurance and will ultimately be the weakest link in the CSLN architecture. A dumb hub would permit a possible intruder to sniff all packets on the network. However, the device does not provide any intelligence for an intruder to subvert the device and hence directly interact with or possibly redirect packets. Although packet sniffing introduces a possible risk, it will be mitigated by the CSLN *trusted channel* which provides packet level confidentiality and integrity. These security elements will provide data protection against a possible intruder gaining intelligence from packet sniffing.

A smart switch would provide the CSLN the ability to stop packet collisions at the networking device which provides for much greater network efficiencies. Smart switching devices also prevent an intruder from sniffing all packets on the network. However, an untrusted smart network switching device presents the possibility for malicious activity (e.g., denial of service and packet redirection) via subversion of the device. Once again, the CSLN *trusted channel* provides packet level confidentiality and integrity and thus provides data protection against a possible intruder gaining intelligence from packet sniffing or redirection.

Both devices have networking and security strengths and weaknesses. Given the networking efficiencies gained by using a smart switch and the risk mitigation provided by invoking the *trusted channel*, the smart switch is recommended for use. However, the

MYSEA certification and accreditation authority will make the ultimate decision based upon the evaluated risk to the architecture.

4. Type I Encryption Device

The CSLN architecture is flexible enough to allow for a TCM that is geographically dislocated from the trusted physical environment of the MYSEA server. However, IPsec Type II cryptographic algorithms do not currently meet DoD Type I requirements for the protection of classified data. Therefore, if a TCM is utilized external to the MYSEA server trusted physical environment, then Type I approved encryption devices are required to cover all communications between the two devices.

B. CONNECTED SINGLE LEVEL NETWORK SOFTWARE MODULES

Modifying MYSEA to incorporate MLS network interfaces requires the development of several custom software modules on both the MYSEA server and TCM. The MYSEA server will require the ability to extract and use an implied *sensitivity level* from each inbound connection and ensure each outbound request destined for a single-level network is at the authorized *sensitivity level*. The MYSEA server and the TCM will require the ability to manage *trusted channel* connections by initiating and receiving PCC protocol requests. The TCM will require the ability to perform network address translation (NAT) [24] on all inbound and outbound packets. The correct implementation of these software modules will permit MYSEA to deliver a distributed MLS network.

Fellows and Weissman's elements for establishing a secure distributed MLS architecture are all enforced in software, of which, accounting for the fragmented TCB requirement is the most difficult to capture. The Common Criteria does not provide a mechanism for evaluating a distributed TSF. Specifically, the Common Criteria does not permit us to evaluate the TSF of the MYSEA server in a context where the *trusted channel* extends its TSF to include the TCM. To overcome this limitation, separate TSFs are developed for the MYSEA server and the TCM to ensure they both enforce separate TSP that permits them to discount the state and delay of the other device while maintaining their own secure state. Separate TSFs mandating each device maintain its own independent secure state also addresses Fellows' fault tolerance requirement for distributed MLS architectures. By ensuring each security device can independently

maintain a secure state, any fault introduced into the system will be systematically handled by both security devices. This functionality ensures that any faults lead to a *fail-secure* state by both devices.

Fellows' *Entelechy* property and the use of trusted protocols to ensure the end-to-end security of the *trusted channel* are additional elements required to deploy a distributed MLS architecture. The *trusted channel* will be based upon the PCC protocol, which will be invoked in software and provide confidentiality. The PCC protocol operates as a trusted process in the IP stack providing end-to-end security for all CSLN connections and will provide the encryption mechanism and key management necessary to enforce Fellows' *Entelechy* property. Wiessman further reinforces the notion that encryption is one of the key elements necessary to establish a *trusted channel*. Specifically, he stipulates that an absolute requirement for protecting a distributed MLS environment is the establishment of a *trusted path (trusted channel)* through the use of a *crypto-seal*.

Fellows' Δ *property* adds one further requirement necessary to create a distributed MLS architecture. The correct and tamperproof administration of the MYSEA server and TCM is imperative if MYSEA is to maintain a secure state. Only authorized security administrators will be allowed to create the initial state necessary to begin MYSEA operations. Likewise, administrative modifications to the MYSEA server and the TCM will only be permitted by authorized security administrators. The TCM will provide no user-level interface to the trusted processes and database necessary to operate CSLN connections.

1. MYSEA Server

The MYSEA server is responsible for the enforcement of all security decisions regarding the correct handling of all data running at multiple *sensitivity levels*. Two new trusted daemons and three new trusted databases are required to implement this functionality: a Trusted Channel Server (TCS); a Secure Connection Server (SCS); a Trusted Channel Database (TCDB), a Secure Connection Inbound Database (SCIDB) and a Secure Connection Outbound Database (SCODB). Additionally, modification to the current IP stack to support the PCC protocol is required to enable the *trusted channel*.

a. *Trusted Channel Server*

The TCS will be responsible for two critical CSLN functions. First the TCS will manage a proprietary MYSEA Security Association protocol necessary to initiate and terminate each *trusted channel* through the PCC protocol. This protocol will be similar in function to the Internet Key Exchange, Version 2 (IKE2) protocol [25], but with a simpler, streamlined implementation. High level design details concerning the MYSEA security association protocol are found in Appendix B of this thesis [26]. Second, the TCS will be responsible for associating an explicit *sensitivity level* to all inbound connections and will be responsible for performing *sensitivity level* equivalence checks on all outbound CSLN connections.

The TCS will receive the implicit sensitivity level of each inbound connection from the PCC protocol handler to derive an explicit *sensitivity level* for that connection. The data and explicit *sensitivity level* will be recorded in the SCIDB for use by the SCS. The TCS will query the SCODB for the explicit *sensitivity level* of each outbound connection and ensure the requested single-level network has an equivalent *sensitivity level*. The TCS will work in concert with the TCDB to conduct these equivalency checks.

b. *Secure Connection Server*

The SCS is a trusted daemon that lies between the TCS and the application protocol servers running on the MYSEA server. The SCS listens on specified ports for all connection attempts and maintains dual responsibilities with respect to the direction of each received connection request. For an inbound connection, the SCS will use information stored in the SCIDB for assigning an explicit *sensitivity level* to each application protocol server it spawns as a result of the incoming connection. For an outbound connection, the SCS extracts the destination IP address of each connection and the *sensitivity level* of the requesting application protocol server and writes the data to the SCODB for use by the TCS.

c. *Trusted Channel Database*

The TCDB is a static database that maintains a record consisting of the IP address of each TCM (the implicit *sensitivity level*), the explicit *sensitivity level* linked to the TCM and the IP address space of the associated single-level networks for which the

TCM is responsible for providing security services. The TCDB provides *read-only* access to the TCS to enable this trusted daemon to determine the explicit *sensitivity level* of each connection.

d. *Secure Connection Inbound Database*

The SCIDB maintains a record consisting of the original source IP address of each inbound connection and the explicit *sensitivity level* of the TCM. The SCIDB is both read and written by the TCS and is *read-only* for the SCS. The TCS receives the IP address of the TCM and the original source IP address of the original packet from the PCC protocol handler on each inbound connection. The TCS queries the TCDB to ensure that the connection is from an authorized TCM and to retrieve the explicit *sensitivity level* of the authorized TCM. The TCS proceeds to write both the explicit *sensitivity level* of the TCM and original source IP address to the SCIDB for later use by the SCS. The TCS is responsible for deleting each record in the SCIDB upon connection termination. The SCS reads the record of each connection from the SCIDB and spawns the requested application protocol server at the explicit *sensitivity level* of the CSLN.

e. *Secure Connection Outbound Database*

The SCODB maintains a record consisting of the destination IP address of each outbound connection and the *sensitivity level* of the originating application process. The SCODB grants the SCS write access to create a record for each outbound connection and read access to check the created record so that all subsequent packets directed toward that connection do not need an additional entry in the database. The SCODB grants the TCS read access to perform a *sensitivity level* equivalence check before permitting the SCS to send the outbound connection request into the networking stack for further processing. The SCODB grants the TCS write access so that a record can be deleted upon termination of the connection.

f. *Protected Communications Channel Protocol handler*

The PCC protocol is the primary enabler necessary to establish a *trusted channel* between the MYSEA server and TCM. As identified in Chapter III and further detailed in Appendix B of this thesis [26], the PCC protocol is an IPsec conformant protocol in accordance with the “Security Architecture for Internet Protocol, Draft-IETF-IPsec_RFC 2401bis” [27]. IPsec provides the basic information assurance (IA) tenets of

authentication, confidentiality, integrity and anti-reply, all of which are necessary to provide *trusted channel* communications. Invoking these tenets of IA permits the MYSEA server to leverage two critical features provided by the *trusted channel*. First, they make possible the ability to ascertain an implied *sensitivity level* for all inbound and outbound MYSEA server communications by establishing an unforgeable, authenticated link between the MYSEA server and TCM. The implied *sensitivity level* of each CSLN is directly linked to the IP address of its TCM. By binding the implicit *sensitivity level* to the IP address of its TCM, the TCS receives the data necessary to authenticate the *sensitivity level* of each connection. Second, they provide a mechanism to enforce data segregation for all connections between the MYSEA server and each of its associated TCMs.

The PCC protocol handler running on the MYSEA server will operate as prescribed by the “RFC 2401bis” in the host system IP stack. However the PCC protocol will implement a custom API in order to extract the source IP address from the inner and outer headers of each packet and subsequently pass that extracted data to the TCS to derive the explicit *sensitivity level* of the connection. The TCS will be required to respond to this request before each packet may proceed.

2. Trusted Channel Module

The TCM provides the basic networking and security functionality required to connect the MYSEA server to a large number of CSLNs. The TCM will provide a secure interface between the MYSEA server and its associated single-level network by establishing a *trusted channel* for all inbound and outbound connections. The TCM will utilize the CISR high assurance TCX kernel to enforce its TSP. Two trusted daemons are required for the TCM to facilitate *trusted channel* services to the MYSEA server. These two daemons are the TCS and the NAT server. Furthermore, the TCM requires the implementation of the PCC protocol to bind the trusted channel between itself and the MYSEA server. The IPsec and NAT functions on the TCM are defined in more detail in Chapter III and in Appendix B of this thesis [26]. Further explanation of the high level design of each daemon is found in those sections.

a. Trusted Channel Server

Similar to the MYSEA server, the TCM will implement a TCS to manage the proprietary MYSEA Security Association protocol necessary to initiate and terminate each trusted channel through the PCC protocol.

b. Network Address Translation Server

The TCM will perform NAT on all inbound and outbound packets. The NAT server will perform destination NAT on all inbound packets and source NAT on all outbound packets. This functionality will allow the MYSEA architecture to not only mask the IP space of its CSLNs, but will also permit MYSEA to comply with current DoD policy, which mandates the separation between unclassified and classified IP space domains (e.g., NIPRNET, SIPRNET, JWICS, etc...). Furthermore, The Defense Information Systems Agency “SIPRNET Classification Guide” [28] prohibits divulging the association of a classified IP address to its name or system to a domain of lower classification. To overcome the networking issues created by this policy, the TCM will perform NAT operations on all packets entering and exiting the CSLN.

c. Protected Communications Channel Protocol

The TCM will implement a “RFC 2401bis” IPsec conformant protocol using the proprietary MYSEA Security Association protocol, which will be discussed in Chapter III and Appendix B of this thesis [26]. No special API will be necessary as the TCM does not perform any *sensitivity level* processing and thus does not need to be aware of its implied *sensitivity level*.

THIS PAGE INTENTIONALLY LEFT BLANK

III. PROTECTED COMMUNICATIONS CHANNEL PROTOCOL

A. OVERVIEW

The MYSEA PCC protocol was first introduced in “A Trusted Connection Framework for Multilevel Secure Local Area Networks” [5]. The PCC protocol was initially designed to establish a *trusted path* between a TPE device and the MYSEA server to support authentication, confidentiality and integrity for all MYSEA MLS LAN communications. Several communications security protocols were analyzed as candidates for the PCC protocol. Based upon several factors which are discussed in Section III A, the IPsec protocol [27] was recommended to be the underlying engine enabling the PCC protocol. Although the PCC is a critical element of the MYSEA architecture, it has yet to be designed and implemented. This document and Appendix B of this thesis [26] are intended to formalize the high level design requirements for the PCC protocol.

The CSLN architecture also requires the use of the PCC protocol. Many elements of the originally recommended PCC protocol remain valid for establishing a *trusted channel* between the MYSEA server and the TCM. As discussed in Chapter II and further defined in Chapter IV, a *trusted channel* also requires authentication, confidentiality and integrity. The PCC protocol design specifications for CSLN operations have the following additional requirements. The PCC protocol will require the use of a security gateway device (i.e., TCM) that can securely interoperate with the MYSEA server. Both the MYSEA server and the TCM will be able to initiate the *trusted channel*. The PCC protocol will require a protocol that can facilitate layer 4 and above communications. The PCC protocol will incorporate anti-replay protection mechanisms to provide protection from maliciously retransmitted packets on the CSLN. These additional PCC protocol requirements reinforce the selection of IPsec as the best candidate protocol for establishing a *trusted channel* within the CSLN architecture. The next section will describe the IPsec protocol and how it is used to support the PCC.

B. IPSEC PROTOCOL

IPsec was selected as the basis for the MYSEA PCC protocol due to its flexibility, scalability and security features, allowing the PCC protocol to meet all *trusted channel* requirements as defined by the Common Criteria. Not only does IPsec provides authentication, confidentiality and integrity mechanisms, but it also includes protection against replay attacks and a security policy database (SPD) that permits/denies all connections based upon a strictly defined set of parameters. Additionally, IPsec is virtually transparent to all protocols above layer 3 [27], which permits IPsec to seamlessly transport all layer 4 and above communications. IPsec uses the SPD in concert with the Security Association database (SAD) to govern all communications flow.

1. Security Policy Database

The IPsec protocol handler controls access to each network interface by making an explicit decision concerning each inbound and outbound packet. The IPsec SPD defines the rules on how each packet is processed. Based upon a set of parameters specified by a security administrator in the SPD, IPsec either discards or permits each packet. For permitted packets, the security parameters contained in the SPD are used to determine whether the packet is processed by the IPsec protection mechanisms or whether it is permitted to bypass those mechanisms and freely pass to the next layer. The SPD maintains an ordered access control list with router-like specificity. This list is used to analyze the attributes of each packet and determine the appropriate action for that packet (e.g., permit with IPsec protection, permit without IPsec protection or discard). If the appropriate action for a packet is to permit with IPsec protection, then information about the cryptographic algorithms designated for use will also be found in the SPD. The IPsec protocol handler will use the data from the SPD to construct a new security association for that connection in the Security Association Database (SAD) – defined in the next section.

The SPD is a trusted database with write access being strictly limited to only the security administrator. Entries in the SPD are based upon the overall security policy for all connections permitted or denied into the host device. The SPD should be configured

before connecting the host device to an untrusted network and should deny all connections that are not expressly permitted. The SPD contains selectors that are equivalent to the access control lists found in a stateless network boundary layer protection device (e.g., router or firewall). At a minimum, the SPD filters on the source and destination IP address and any designated data found at or above the IP layer (e.g., TCP, UDP, ICMP, ports, etc...).

The integrity of the SPD and SAD is critical to the secure operation of each IPsec connection. The authority to administer the IPsec SPD will be strictly controlled by the security administrator in compliance with the Δ *property*.

2. Security Association Database

The SAD is a dynamic database that stores the parameters necessary for the IPsec protocol handler to apply the cryptographic algorithms on an active IPsec connection. A SAD record is created by the IPsec Security Association (SA) protocol, later discussed in Chapter III.B.5. Upon the initial set-up of an IPsec connection, the SA protocol handler negotiates the values necessary to enable the authentication, encryption and integrity security services afforded by IPsec. Once negotiated, these values are stored in the SAD so that all further IP packets traversing the connection can utilize the existing SA and avoid the overhead of additional SA negotiation.

3. Modes of Operation

IPsec can operate in two distinct modes of operation: 1) Transport Mode; and 2) Tunnel Mode. Transport Mode is generally implemented when connections are between two end point devices. The Transport Mode header retains the header of the original IP packet and therefore provides no protection against basic traffic analysis. Figure 6 illustrates the packet transformation in transport mode using the Encapsulating Security Protocol to provide both confidentiality and integrity protection on the packet. Tunnel Mode is implemented when at least one of the security devices is a security gateway. Tunnel Mode affords the security gateway the ability to provide additional traffic flow protection by masking the original IP header through encryption. A Tunnel Mode IPsec Encapsulating Security Protocol (ESP) packet is illustrated in Figure 7 which implements both confidentiality and integrity – ESP will be explained in Chapter III.B.4. MYSEA

will utilize Tunnel Mode to maximize security protection mechanisms between the MYSEA server and its associated TCMs. Tunnel Mode will permit the CSLN to provide full protection for every original IP packet flowing through the CSLN.




Figure 6. Transport Mode IPsec Packet using ESP Protocol




Figure 7. Tunnel Mode ESP IPsec Packet

4. IPsec Security Protocols

MYSEA will incorporate the IPsec security services necessary to establish a *trusted channel*. IPsec provides its cryptographic security services through the implementation of the Encapsulating Security Protocol (ESP) and the Authenticating Header (AH) protocol. The Encapsulating Security Payload (ESP) Protocol [29] provides authentication, confidentiality and integrity for an IP packet. However, ESP does not provide complete integrity protection for an IPsec packet because it leaves the original IP header unprotected in Transport Mode and the new IP header unprotected in Tunnel Mode as illustrated in Figures 6 and 7. If either the Transport or Tunnel Mode IP header is left unprotected, malicious actions (e.g., IP spoofing, redirects, etc...) against the header may adversely impact the integrity of the *trusted channel*. This deficiency is unacceptable in the CSLN architecture as the implicit sensitivity level of each packet is directly tied to the TCM's IP address, which is located in the new IP header. The integrity protection deficiency inherent to ESP can be overcome by encapsulating the ESP packet in the Authenticated Header (AH) Protocol [30]. AH provides packet authentication and integrity to every field in the IPsec packet, less some mutable values

found in the IP header (e.g., TTL value), as illustrated in Figure 8. The combination of these two security protocols is required for MYSEA to provide a *trusted channel* and protect against malicious actions affecting the authentication, confidentiality, or integrity of CSLN communications. .




Figure 8. AH/ESP Tunnel Mode Protected Packet

5. Security Association and Key Management

IPsec provides two methods for negotiating a Security Association (SA) and managing cryptographic keys. The simplest method is a manual SA configured with static SA management variables and cryptographic keys. This method typically incorporates the use of pre-selected symmetric keys to be used in conjunction with its cryptographic algorithms. Although the manual method presents the easiest solution for establishing an IPsec security association, its key management architecture lacks the ability to scale with large, highly distributed network architectures. Additionally, IPsec loses its ability to provide anti-replay protection with the use of manual security associations as it does not allow for a connection by connection negotiation of a unique counter necessary to prevent replay attacks.

Automated SA and key management provides the second method for negotiating new security associations and managing cryptographic keys. Automated SA management provides IPsec the capability to scale with large network implementations, to negotiate its anti-replay protection mechanisms and to create several distinct SAs between the same two hardware devices. The IKE2 protocol [25] is the standard for IPsec automated SA negotiation and key management. However, IKE2 introduces unneeded complexity and negotiation overhead into MYSEA. As a result, MYSEA will use a proprietary automated SA and key management algorithm to bootstrap each IPsec connection. This algorithm will use a simplified, custom key exchange protocol that will

take advantage of the master/slave relationship between the MYSEA server and TCM. The MYSEA SA protocol will use public key cryptography digital certificates to secure the initial SA setup. The MYSEA server alone will be responsible for key generation and distribution to its associated TCMs on a connection-by-connection basis. The reason for implementing a proprietary automated SA and key management algorithm is two-fold. First, this protocol will reduce the IKE2 connection negotiation overhead. Second, this protocol will reduce the size and complexity of the SA mechanism. This will contribute to its understandability, thus making it appropriate for a high assurance implementation. Details concerning the automated key exchange protocol for MYSEA are found in Appendix B of this thesis [26].

6. IPsec Cryptographic Algorithms

Part of the rich set of security services delivered by IPsec is the choice of implementing one of several cryptographic confidentiality and integrity algorithms. Nevertheless, MYSEA will streamline the security association process by only using the strongest cryptographic algorithms available for IPsec integration. Currently, the National Institute for Technology (NIST) standard for providing confidentiality is the Advanced Encryption Standard (AES-128-CBC) algorithm [31]. The NIST standard for integrity is the Secure Hash Algorithm (SHA-1) [32]. Applying these two algorithms together through the ESP and AH security protocols provides IPsec the ability to facilitate authentication, confidentiality and integrity for all packets transiting the protected network interface. In the future, MYSEA will continue to adopt the strongest cryptographic algorithms available for integration. If for any reason, one or both of these algorithms are found to be vulnerable to a cryptographic attack, then MYSEA will quickly move to adopt an accepted replacement algorithm.

7. IPsec Placement

The flexible nature of IPsec allows for designers to choose from one of three possible implementations. The most advantageous implementation is the embedment of IPsec code into the native operating system IP stack. This implementation permits IPsec

to operate more efficiently. However, native integration of IPsec into the existing operating system requires the ability to access and modify the existing source code from the IP stack.

A “bump-in-the-stack” (BITS) design is an alternative method for integrating IPsec into an operating systems IP stack. BITS permits designers to integrate IPsec into the local networking stack when access to the operating source code is unavailable. BITS inserts IPsec into the local networking stack between the IP layer and the local network drivers in the data link layer.

A “bump-in-the-wire” (BITW) design is used when an external device enables the IPsec functionality. Typically, a BITW design refers to the use of a separate security gateway device (e.g., router or firewall) with its own IP address to implement IPsec. A BITW design may also be used internal to a host-based system through the use of a custom IPsec cryptographic module.

The CSLN requires the integration of IPsec in both the MYSEA server and TCM. By definition, the TCM will use the BITW method for delivering IPsec services, but more specifically, IPsec will be natively integrated into the operating system IP stack of the TCM. If possible, the MYSEA server should use a *native* IPsec integration as well. If a native integration proves impossible due to a lack of access to the source code of the operating system, then the MYSEA server may use a “bump-in-the-stack” (BITS) IPsec integration.

C. RESIDUAL RISK

The IPsec protocol provides authentication, confidentiality and integrity, which are all required for creating the *trusted channel* necessary to enable CSLN operations. Nonetheless, some residual risk remains in the architecture. First, although IPsec tunnel mode hides the IP address of the originating host, basic traffic analysis can still occur against the packets transiting the CSLN. Heavy traffic transiting to one single-level network or another may provide insight into special intelligence events or operations. It should be noted that the DoD has already accepted this risk with its implementation of the current generation of IP encryptors. Second, the CSLN may permit the creation of covert channels by manipulating the mutable fields of the IP header. This vulnerability

could lead to the leakage of classified data into the CSLN. Finally, IPsec implements Type II cryptographic algorithms. Current DoD policy requires the use of Type I algorithms to protect classified data [33].

Although mitigation techniques exist to reduce the risk associated with these vulnerabilities, future work for MYSEA architects should include a detailed vulnerability analysis and determination of the best method for mitigating the risk. This analysis should be followed with a reasoned risk management decision to accept or not accept the remaining residual risk, if any.

D. SUMMARY

The CSLN architecture requires the implementation of a PCC protocol in order to provide a *trusted channel* as delineated by the Common Criteria. The IPsec protocol satisfies all *trusted channel* requirements and is chosen as the basis for the PCC protocol. By combining the IPsec ESP and AH security protocols, the PCC will provide authentication, confidentiality and integrity to every packet traversing the CSLN. The use of tunnel mode will permit the attachment of an implicit *sensitivity level* to these packets by directly associating the IP address of the TCM to the connection. The proprietary security association and key management protocol will afford MYSEA a streamlined and provable method for bootstrapping each IPsec connection. The culmination of the IPsec security functionality will afford data separation between each packet arriving and departing the MYSEA server at multiple *sensitivity levels*. Furthermore, this security functionality will enable the MYSEA server to provide an MLS network interface, which in turn creates the capability to provide a truly distributed MLS network architecture.

IV. THE MYSEA TRUSTED CHANNEL MODULE

Development of the TCM will be in accordance with the Common Criteria. However, complying with the full spectrum of requirements necessary to generate a formal Security Target is beyond the scope of this thesis. This chapter only provides a general description of the threats, objectives and assumptions necessary to design the TCM. The term Target of Evaluation (TOE) refers to the entire TCM, which may be composed of several components, some of which may have their own TOE.

As previously discussed in Chapter I, three documents provided the guidance for developing this threats and objectives analysis: The “Consistency Instruction Manual for Development of U.S. Government Protection Profiles for use in Medium Robustness Environments” [19]; the “ST for Cisco IOS/IPsec” [20] ; and the “U.S. Government Protection Profile for Separation Kernels in Environments Requiring High Robustness – Version 0.621” [21].

A. TOE SECURITY ENVIRONMENT

1. Assumptions

The following table list the assumptions made concerning the TOE environment.

| | |
|---------------------------|--|
| A.PHYSICAL [21] | It is assumed that the IT environment provides the TOE with appropriate physical security commensurate with the value of the IT assets protected by the TOE. |
| A.TRUSTED_INDIVIDUAL [21] | If an individual is allowed to perform procedures upon which the security of the TOE may depend, it is assumed that the individual is trusted with assurance commensurate with the value of the IT assets. |
| A.TRAINING [20] | As the security functions of the TOE can be compromised due to errors or omissions in the administration of the security features of the TOE, it is assumed that administrators of the TOE have been trained to enable them to securely configure the TOE. |

| | |
|-------------------|--|
| A.TRUSTED-CA [20] | As the security functions of the TOE when configured to use digital certificates can be comprised if the Certificate Authority (CA) that issued the certificates is not operated in a trusted manner, it is assumed that if the TOE is configured to use digital certificates, the issuing CA is trusted or evaluated to at least the same level as the TOE. |
|-------------------|--|

Table 1. Security Usage Assumptions

2. Threats

The following table provides a sketch of the threats anticipated against the TOE environment.

| | |
|--------------------------|--|
| T.ADMIN_ERROR [21] | An administrator may incorrectly install or configure the TOE (including the misapplication of the principle of least privilege to limit the damage that can result from accident, error, or unauthorized use), or install a corrupted TOE resulting in ineffective security mechanisms. |
| T.ALTERED_DELIVERY [21] | The TOE may be corrupted or otherwise modified during delivery such that the on-site version does not match the master distribution version. |
| T.BAD_RECOVERY [21] | The TOE may be placed in an insecure state as a result of unsuccessful recovery from a system failure or discontinuity. |
| T.COVERT_CHANNEL_EXPLOIT | Unauthorized data may be tunneled through the TOE. |
| T.CRYPTO_COMPROMISE [21] | A malicious subject may cause key, data or executable code associated with the cryptographic functionality to be inappropriately accessed (viewed, modified, or deleted), thus compromising the cryptographic mechanisms and the data protected by those mechanisms. |
| T.INSECURE_STATE [21] | When the TOE is initially started or restarted after a failure, the security state of the TOE may be in an insecure state. |

| | |
|---------------------------------|--|
| T.POOR_DESIGN [21] | Unintentional or intentional errors in requirements specification or design of the TOE may occur, leading to flaws that may be exploited by a malicious subject. |
| T.POOR_IMPLEMENTATION [21] | Unintentional or intentional errors in implementation of the TOE design may occur, leading to flaws that may be exploited by a malicious subject. |
| T.POOR_TEST [21] | Lack of or insufficient tests to demonstrate that all TOE security functions operate correctly (including in a fielded TOE) may result in incorrect TOE behavior being undiscovered. |
| T.ATTACK [20] | An attacker (whether an insider or outsider) may gain access to the TOE and compromise its security functions by altering its configuration. |
| T.AUDIT_COMPROMISE [19] | A malicious user or process may view audit records, cause audit records to become or modified, or prevent future audit records from being recorded, thus masking a user's action |
| T.REPLAY [19] | An IT entity may gain inappropriate access to unauthorized data traversing the CSLN by replaying IP packets through the network. |
| T.RESOURCE_EXHAUSTION [19] | A malicious IT entity may block access to the <i>trusted channel</i> by exhausting the resources on the TOE required to initiate a new connection. |
| T.SPOOFING [19] | A malicious IT entity may misrepresent itself as the TOE to obtain unauthorized data |
| T.MALICIOUS_TSF_COMPROMISE [19] | A malicious unauthorized IT entity may cause TSF data or executable code to be inappropriately accessed (viewed, modified or deleted) |

Table 2. Anticipated Threats Against TOE

3. Organizational Security Policies

The following table lists some of the organizational security policies applicable to the TOE.

| | |
|--|---|
| P.ADMIN_ACCESS [19] | Administrators shall be able to administer the TOE locally through the protected communications channels |
| P.ACCOUNTABILITY [21] | The TOE shall provide the capability to make available information regarding the occurrence of security relevant events. |
| P.CRYPTOGRAPHY [21] | The TOE shall use NIST FIPS validated cryptography as a baseline with additional NSA-approved methods for key management (i.e., generation, access, distribution, destruction, handling, and storage of keys) and for cryptographic operations (i.e., encryption, decryption, signature, hashing, key exchange, and random number generation services). |
| P.LEAST_PRIVILEGE [21] | The TOE shall be designed such that the principle of least privilege is applied to limit the damage that can result from accident, error or unauthorized use. |
| P.RATINGS_MAINTENANCE [21] | A plan for procedures and processes to maintain the TOE's rating must be in place to maintain the TOE's rating once it is evaluated. |
| P.SYSTEM_INTEGRITY [21] | The TOE shall provide the ability to periodically validate its correct operation and, with the help of administrators if necessary, it must be able to recover from any errors that are detected. |
| P.ADMIN_GUIDANCE [21] | The TOE shall provide documentation and training regarding the correct use of the TOE security features. |
| P.VULNERABILITY_ANALYSIS_AND_TEST [21] | The TOE must undergo independent vulnerability analysis and penetration testing to demonstrate that the TOE is resistant to an attacker possessing a high attack potential. |

Table 3. Organizational Security Policies Applicable to the TOE

B. SECURITY OBJECTIVES

1. Security Objectives for the TOE

The following security objectives will enable the TOE to counter known threats and comply with identified organizational security policies and assumptions.

| | |
|---------------------------------|--|
| O.ACCESS [21] | The TOE will ensure that subjects gain only authorized access to resources that it controls. |
| O.ADMIN_GUIDANCE [21] | The TOE will provide administrators with the necessary information for secure management of the TOE. |
| O.AUDIT_GENERATION [21] | The TOE will provide the capability to detect and generate audit records for security relevant auditable events. |
| O.AUDIT_PROTECTION [19] | The TOE will provide the capability to protect audit information. |
| O.CHANGE_MANAGEMENT [21] | The configuration of, and all changes to, the TOE and its development evidence will be analyzed, tracked, and controlled throughout the TOE's development. |
| O.CORRECT_BOOT [21] | The TOE will provide mechanisms to correctly transfer the TSF implementation and TSF data into the TSF's execution domain. |
| O.CORRECT_TSF_OPERATION [21] | The TOE will provide a capability to test the TSF to ensure the correct operation of the TSF during normal operation. |
| O.COVERT_CHANNEL_ANALYSIS [21] | The TOE will undergo appropriate covert channel analysis to demonstrate that the TOE meets its functional requirement. |
| O.CRYPTOGRAPHIC_PROTECTION [21] | The TOE will support separation of the cryptography from the rest of the TSF. |
| O.CRYPTOGRAPHIC_SERVICES [21] | The TOE will use cryptographic mechanisms to protect the integrity of TOE code and data as it resides within the system and when it is transmitted to other systems. The TOE will also use cryptographic mechanisms to verify the integrity of the TSF code and configuration data during initialization. The cryptographic mechanism will use NIST FIPS validated cryptography as a |

| | |
|---------------------------------|---|
| | baseline with additional NSA-approved methods. |
| O.DISPLAY_BANNER [19] | The TOE will display an advisory warning regarding use of the TOE. |
| O.DOCUMENT_KEY_LEAKAGE [19] | The bandwidth of channels that can be used to compromise key materials shall be documented. |
| O.FUNCTIONAL_TESTING [21] | The TOE will undergo independent security functional testing that demonstrates the TSF satisfies the security functional requirements. |
| O.INSTALL_GUIDANCE [21] | The TOE will be delivered with the appropriate installation guidance to establish and maintain TOE security. |
| O.INTERNAL_LEAST_PRIVILEGE [21] | The entire TSF will be structured to achieve the principle of least privilege among TSF modules. |
| O.MAINTENANCE_MODE [19] | The TOE shall provide a mode from which recovery or initial startup procedures can be performed. |
| O.MANAGE [21] | The TOE will provide all the functions necessary to support the administrative users and authorized subjects in their management of the configuration data, and restrict these functions from use by unauthorized subjects. |
| O.PROTECT [21] | The TOE will provide mechanisms to protect services and exported resources. |
| O.RATINGS_MAINTENANCE [21] | Procedures and processes to maintain the TOE's rating will be documented. |
| O.RECOVERY [21] | Procedures and/or mechanisms will be provided to assure that recovery, such as from system failure or discontinuity, is obtained without a protection compromise. |
| O.REFERENCE_MONITOR [21] | The TOE will maintain a domain for its own execution that protects itself and its resources from external interference, tampering, or unauthorized disclosure. |
| O.REPLAY_DETECTION [19] | The TOE will provide a means to detect and reject the replay of TSF data traversing the CSLN. |
| O.RESIDUAL_INFORMATION [21] | The TOE will ensure that any information contained in a protected resource is not released when the resource is reallocated. |

| | |
|------------------------------------|---|
| O.RESOURCE_SHARING [21] | The TOE will provide mechanisms that mitigate attempts to exhaust TOE resources (e.g., system memory and processing time). |
| O.ROBUST_TOE_ACCESS [19] | The TOE will provide mechanisms that control a user's logical access to the TOE and to explicitly deny access to specific users when appropriate. |
| O.SECURE_STATE [21] | The TOE will provide mechanisms to transition the TSF to a secure state during start-up. The TSF will be designed to maintain a secure state. |
| O.SELF_PROTECTION [19] | The TSF will maintain a domain for its own execution that protects itself and its resources from external interference, tampering or unauthorized disclosure. |
| O.SOUND_DESIGN [21] | The TOE will be designed using sound design principles and techniques. The TOE design, design principles and design techniques will be adequately and accurately documented. |
| O.SOUND_IMPLEMENTATION [21] | The implementation of the TOE will be an accurate instantiation of its design. |
| O.THOROUGH_FUNCTIONAL_TESTING [19] | The TOE will undergo appropriate security functional testing that demonstrates the TSF satisfies the security functional requirements. |
| O.TRUSTED_DELIVERY [21] | The integrity of the TOE must be protected during the initial delivery and subsequent updates, and verified to ensure that the on-site version matches the master distribution version. |
| O.TSF_INTEGRITY [21] | The TOE will be able to verify the integrity of the TSF code and data. |
| O.USER_GUIDANCE [21] | The TOE will provide users with the necessary information for secure use of the TOE. |
| O.VULNERABILITY_ANALYSIS_TEST [21] | The TOE will undergo independent vulnerability analysis and penetration testing to demonstrate the design and implementation of the TOE does not allow attackers with high attack potential to violate the TOE's security policies. |

Table 4. TOE Security Objectives

2. Security Objectives for the Environment

The following table specifies the security objectives for the environment, which are not directly countered by the TOE security objectives or by the organizational security policies and assumptions.

| | |
|----------------------------|--|
| OE.PHYSICAL [21] | Physical security will be provided for the TOE by the IT environment commensurate with the value of the IT assets protected by the TOE. |
| OE.TRUSTED_INDIVIDUAL [21] | If an individual is allowed to perform procedures upon which the security of the TOE may depend, that individual is trusted with assurance commensurate with the value of the IT assets. |

Table 5. Security Objectives for the Environment

V. FUTURE WORK AND CONCLUSION

A. FUTURE WORK

MYSEA is an ambitious project with several ongoing research efforts. The research presented in this document provides the focus for several future research initiatives that will serve to enhance both the security and capabilities of MYSEA. These research initiatives were considered beyond the scope of this document.

1. Formal TCM Security Target

This document provides the initial high-level functional requirements for the TCM. The next logical step is the detailed design of the TCM using a formal Common Criteria Security Target methodology. Furthermore, this design should identify the form factor requirements necessary to host the TCM

2. TCM Fail-Over

To ensure a robust CSLN design, a fail-over mechanism should be incorporated into future versions of the TCM. This fail-over design could emulate the Hot Standby Routing Protocol (HSRP) currently fielded by Cisco Systems Inc. This functionality will ensure the failure of a single TCM does not cripple connectivity to a critical single-level network [34].

3. Single-Level Network User Identification and Authentication

The current CSLN architecture only permits identification and authentication of the single-level network to the level of the corresponding CSLN. As a result, individual users accessing the MYSEA server from the single-level network cannot currently be authenticated. Sound security policy dictates that all interactions with the MYSEA server should be associated with a *sensitivity level* and an authenticated user. Future versions of the CSLN architecture should include a mechanism for identifying each individual user accessing the MYSEA server from a single-level network through the CSLN.

4. MYSEA Security Association Protocol

MYSEA intends to design and develop a proprietary Security Association Protocol to perform the initial binding of all PCC connections. Although a high level

design of that protocol is presented in this document, future work should include a thorough security analysis of the protocol followed by its low level design and development.

5. MLS IP Encryptors

Implementing a *trusted channel* at the physical layer would not require the development of a PCC protocol or a TCM. It would require the addition of separate, external gateway security devices located on the physical wire in front of the MYSEA server and its associated single-level networks. Several commercial and government virtual private network security devices are strong candidates to implement a Layer 1 solution. National Security Agency approved Type I IP layer encryption devices currently exist and would satisfy all security requirements necessary to implement a trusted channel. The use of these Type I certified devices would mitigate the residual risk inherent in utilizing Type II approved devices. Additionally, the use of Type I encryption devices would allow the MYSEA server to directly connect to any geographically distant packet switched single-level network.

Two major problems currently prevent the implementation of a Layer 1 security solution. First, the present generation of Type I IP encryptors is not certified to handle MLS data streams. Certification to handle an MLS data stream is a requirement for any Type I encryptor processing MYSEA server communications. Second, the Type I encryptor would lie outside the TSF of the MYSEA server and thus provide no trusted mechanism to associate an explicit *sensitivity level* to the incoming connection. An explicit *sensitivity level* could be extracted from the implied *sensitivity level* associated with the original source IP address of the packet, but this implementation could be compromised by IP spoofing the original source IP address inside the untrusted single-level network. With no mechanism to authenticate that the packet was not compromised by IP spoofing, a Layer 1 solution using existing devices would leave the CSLN architecture with considerable residual risk.

Nonetheless, MLS Type I IP encryptors are a long-term requirement for the DoD. *Future work* for MYSEA should include reassessing the CSLN architecture upon the certification of these devices. Depending upon the features incorporated into an MLS

Type I encryptor, MYSEA may have the opportunity to streamline its CSLN architecture and increase its security posture by using this device and its associated cryptographic algorithms.

6. MLS Server-to-MLS Server Connectivity

This document provides an architecture for connecting an MLS server to a large number of single-level networks. In the future, MYSEA will transition to a distributed MLS network, which relies upon several MLS servers each of which is able to manage data at different *sensitivity levels*. To meet this goal, MYSEA requires an extension to the PCC protocol to support sharing of data at different *sensitivity levels* securely between MYSEA servers.

7. Protected Communications Channel Residual Risk Analysis

As discussed in Chapter III C, the PCC protocol cannot mitigate all risk associated with the use of a *trusted channel*. Further risk analysis should be conducted on the possibility of exploiting the PCC through the use of basic traffic analysis, covert channels, and the execution of Type II cryptographic algorithms. All of these potential vulnerabilities have known risk mitigation techniques that should be studied and applied to the CSLN architecture. Through this analysis, an intelligent risk management decision can be made concerning any residual risk associated to the CSLN architecture.

B. CONCLUSION

This document proposes an extension to the current Monterey Security architecture to support simultaneous connection management and protection in a distributed MLS environment. Two previous works were instrumental to understanding the design requirements for this research. First, Wilson's, "A Trusted Connection Framework for Multilevel Secure Local Area Networks" [5] provided the underlying knowledge critical to understanding how the MYSEA server provides its security services. Second, Fellows' paper entitled "The Architecture of a Distributed Trusted Computing Base" [16] provided the fundamental knowledge required to create a distributed MLS architecture.

The ultimate goal of MYSEA is to field a confederation of MLS servers capable of sending and receiving multiple MLS data streams. However, the current DoD

architecture is tied to a substantial number of system high, dedicated and compartmented single-level networks that will remain operational until a robust, provable MLS architecture is fielded. This document describes an architecture that serves to provide a bridge between the current multiple single-level network architectures and the MLS network architecture of the future. Contemporary MLS server technologies are restricted in their ability to connect to a large number of single-level networks by their limited number of physical network interfaces. This constraint prevents the situational awareness gained by fusing data from a large number of single-level networks on a central MLS server. The architecture presented in this document permits the MYSEA server to overcome its limited number of dedicated network interfaces by providing an MLS network interface capable of supporting a large number of single-level networks.

This architecture is enabled by modifying the existing MYSEA server, creating the TCM and by providing a *trusted channel* between these two security critical devices. The *trusted channel* is the key component for creating a distributed TSF between the MYSEA server and TCM. The PCC protocol, based on the IPsec protocol, is responsible for creating the *trusted channel*. The PCC protocol binds the security tenets of authentication, confidentiality and integrity to all CSLN communications. The security mechanisms provided by the *trusted channel* permit the MYSEA server to extract an implied sensitivity level from each connection by positively authenticating each associated TCM to the MYSEA server.

The MYSEA server requires the addition of several trusted daemons and databases to enforce the CSLN TSP. These daemons and databases ensure that the implicit *sensitivity level* of all inbound *trusted channel* communications is bound to a known explicit *sensitivity level* before access to services on the MYSEA server is authorized. Furthermore, these daemons and databases ensure that all outbound communications are permitted only to a CSLN with a *sensitivity level* equivalent to that of the originating MYSEA server process.

The TCM functions as a security gateway by bridging all communication between its associated single-level network and the MYSEA server. The TCM serves as a trusted end-point for all *trusted channel* communications with the MYSEA server. This

functionality permits the MYSEA server to positively authenticate each TCM and as a result, associates an implicit *sensitivity level* with all communications traversing the TCM.

The MYSEA server, TCM and MYSEA *trusted channel* work in concert to support a distributed MLS environment. These components enforce all connection management and *trusted channel* security requirements necessary for the MYSEA server to provide an MLS network interface capable of supporting a large number of single-level networks. The integration of this functionality into MYSEA will enable a truly distributed MLS architecture capable of linking current DoD single-level networks to a repository of information at different *sensitivity levels*.

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX A

CONNECTED SINGLE LEVEL NETWORK SYSTEM REQUIREMENTS DOCUMENT

TABLE OF CONTENTS

| | | |
|--|---|-----------|
| 1. | INTRODUCTION..... | 47 |
| 1.1 | PURPOSE..... | 47 |
| 1.2 | SCOPE | 47 |
| 2. | SYSTEMS OVERVIEW | 49 |
| 2.1 | MYSEA DISTRIBUTED MLS ENVIRONMENT OVERVIEW | 49 |
| 2.2 | CSLN DESCRIPTION | 51 |
| 2.3 | CSLN COMPONENT DESCRIPTIONS | 52 |
| 2.3.1 | MYSEA Server Target of Evaluation Security Function..... | 53 |
| 2.3.1.1 | <i>Trusted Channel Server.....</i> | <i>54</i> |
| 2.3.1.2 | <i>Secure Connection Server</i> | <i>54</i> |
| 2.3.1.3 | <i>Trusted Channel Database</i> | <i>55</i> |
| 2.3.1.4 | <i>Secure Connection Inbound Database.....</i> | <i>55</i> |
| 2.3.1.5 | <i>Secure Connection Outbound Database</i> | <i>56</i> |
| 2.3.1.6 | <i>Protected Communications Channel Protocol Handler.....</i> | <i>56</i> |
| 2.3.2 | Trusted Channel Module Target of Evaluation Security Function Services | 57 |
| 2.3.2.1 | <i>Trusted Channel Server.....</i> | <i>57</i> |
| 2.3.2.2 | <i>Network Address Translation Server</i> | <i>58</i> |
| 2.3.2.3 | <i>Protected Communications Channel Protocol Handler.....</i> | <i>58</i> |
| 2.3.3 | Single Level Networks | 58 |
| 3. | SYSTEM REQUIREMENTS | 59 |
| 3.1 | CONNECTED SINGLE LEVEL NETWORK REQUIREMENTS | 59 |
| 3.2 | MYSEA SERVER REQUIREMENTS | 59 |
| 3.3 | TRUSTED CHANNEL MODULE REQUIREMENTS..... | 60 |
| 3.4 | MYSEA CONNECTION PROTOCOL REQUIREMENTS..... | 61 |
| 3.5 | CONNECTED SINGLE LEVEL NETWORK APPLICATION PROTOCOL SERVICES REQUIREMENTS..... | 61 |
| APPENDIXES..... | | 63 |
| APPENDIX A ABBREVIATIONS, ACRONYMS, AND DEFINITIONS..... | | 63 |
| APPENDIX B REFERENCES | | 65 |

THIS PAGE INTENTIONALLY LEFT BLANK

1. INTRODUCTION

1.1 PURPOSE

Extending the initial design requirements for the Monterey Security Architecture (MYSEA) is the purpose of this systems requirements document. These requirements were initially identified in “A Trusted Connection Framework for Multilevel Secure Local Area Networks” [1]. This document conforms to the original Systems Requirements Document outline, where possible, and modifies the document to account for dissimilar requirements when necessary. Common Criteria standards and terminology will be used to describe all hardware, software, firmware, and their interactions [2].

1.2 SCOPE

This document delineates new MYSEA requirements necessary to extend the functionality of the MYSEA server to allow simultaneous protected access to a large number of single level networks at different *sensitivity levels* while providing simultaneous management of all connections. Some of these networks will be multiplexed through a single multilevel security (MLS) network interface on the MYSEA server. The creation of a Trusted Channel Module (TCM) security device capable of establishing a *trusted channel* component is the key enabler for this functionality. This document establishes the requirements essential to launch a *trusted channel*. Satisfying these requirements is crucial if MYSEA is to provide a robust distributed MLS environment.

THIS PAGE INTENTIONALLY LEFT BLANK

2. SYSTEMS OVERVIEW

2.1 MYSEA DISTRIBUTED MLS ENVIRONMENT OVERVIEW

MYSEA is intended to demonstrate the feasibility of a certifiable MLS architecture. This architecture will provide the end user with commercial-off-the-shelf form, function and features while enforcing high assurance security services with a minimum number of high assurance security components. This architecture is envisioned to consist of a federation of MLS servers with each server supporting a MLS local area network (LAN) and a large number of *Single-Level Networks* (e.g., NIPRNET, SIPRNET, JWICS). Currently, the MYSEA prototype comprises one main segment featuring a MLS LAN. The MLS LAN encompasses the MYSEA server and a set of untrusted thin clients, each having a dedicated *Trusted Path Extension* (TPE) device. The TPE provides a *trusted path* interface to the MYSEA server, enabling the user to utilize the untrusted client to access server data and services at any authorized *sensitivity level*. This functionality distributes the Target of Evaluation (TOE) Security Function (TSF) of the MYSEA server to the untrusted client and is the key enabler for the MLS LAN environment. The MYSEA server also includes a limited number of network interfaces to support a small number of dedicated single-level networks. Each single-level network interfaces with the MYSEA server via a dedicated network interface with a pre-defined *sensitivity level* that corresponds to the *sensitivity level* of the single-level network. This system requirements document seeks to further extend MYSEA functionality by modifying the architecture to support a MLS network interface capable of handling a large number of single-level networks at multiple *sensitivity levels*. This modification will be known as the Connected Single Level Network (CSLN) architecture and is depicted in Figure 1.

One significant modification has occurred to MYSEA since the release of its initial system requirements document. MYSEA upgraded its primary high assurance component to the DigitalNet XTS-400 MLS server running the Secure Trusted Operating Program (STOP) 6.1 [3]. The XTS-400 continues to enforce a MLS policy based upon the Bell and LaPadula [4], and Biba [5] security and integrity models and is

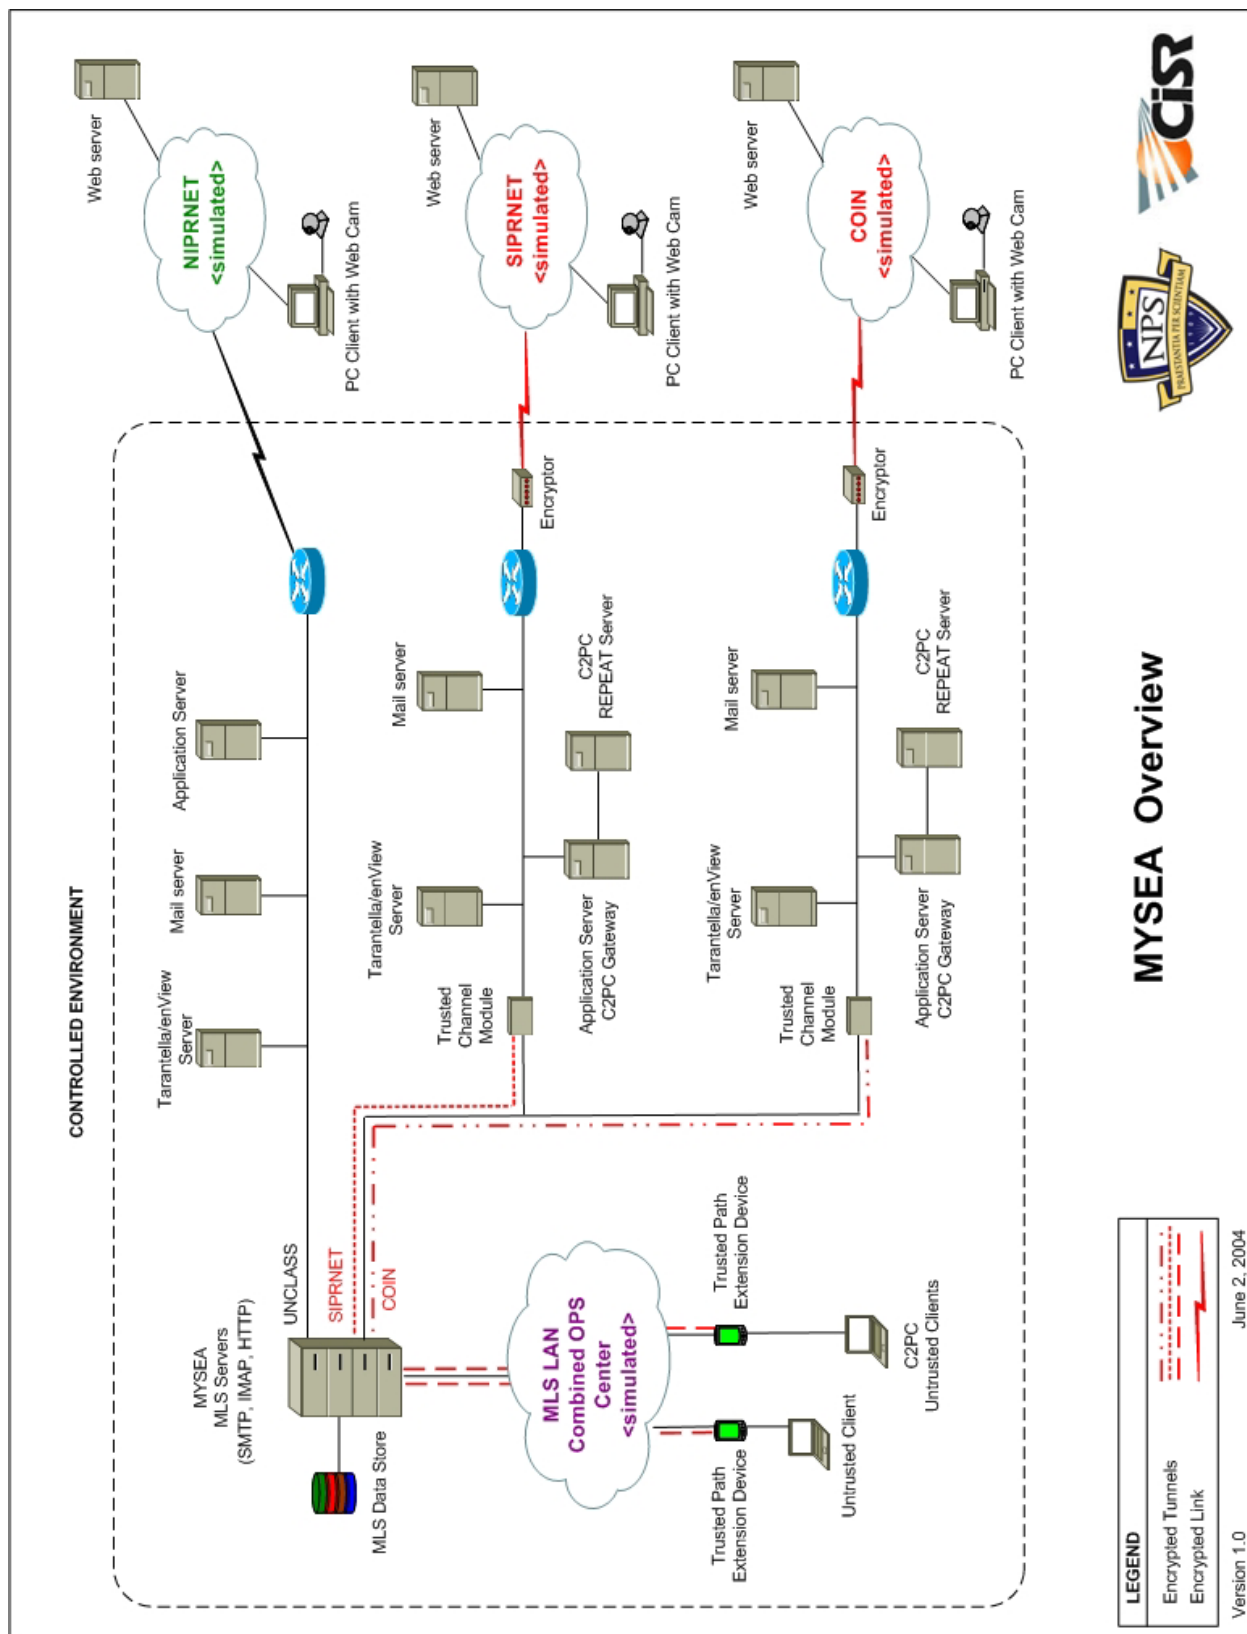


Figure 1. MYSEA with TCM Integration

completing evaluation under the Common Criteria at EAL 5+. Like the XTS-300, the XTS-400 also takes advantage of the Intel x86 chipset four *domains of isolation* architecture to help enforce “*the principle of least privilege*” at the hardware level [6].

Although many features from the XTS-300 were migrated to the XTS-400, one major modification was the level of access given to MYSEA programmers to interface with the STOP. Whereas the XTS-300 allowed programmers “to place a trusted daemon process in the Operating System Services (OSS) domain of operations,” the XTS-400 no longer allows that level of interaction. Hence, trusted MYSEA daemons now operate in the *Application Domain* as shown in Figure 2.

To distribute the TSF of the MLS server to each CSLN routed through a MLS network interface, a TCM is required. The TCM is responsible for creating a *logically isolated and unmistakably distinguishable [1] trusted channel* between itself and the MYSEA server. The *trusted channel* will provide a means to attach an implied *sensitivity level* to all communications between the MYSEA server and the TCM. The TCM will provide the interface for all untrusted CSLN communications addressed to an MLS network interface on the MYSEA server. The TCM, in collaboration with the MYSEA server, will provide a verifiable communications channel between the TCM and MYSEA server. The *trusted channel* will be enabled by the development of a *Protected Communications Channel (PCC)* protocol providing the cornerstone of all *trusted channel* communications. The *trusted channel*, in concert with the MYSEA server and its TCMs, will provide for the authentication, integrity and confidentiality of all CSLN communications.

2.2 CSLN DESCRIPTION

The CSLN architecture provides peer-to-peer connectivity between the MYSEA server and TCM for existing system high, dedicated or compartmented networks. The CSLN is comprised of the MYSEA server, TCM and an untrusted networking device. At all times, CSLN connectivity will be constrained by the *TOE Security Policy (TSP)* [2] enforced by the MYSEA server. MYSEA LAN users will maintain the ability to access applications on various single-level networks protected by CSLN security components operating at the authorized *sensitivity level*. Conversely, users on each single-level

network will maintain the ability to access applications running on the MYSEA server at the authorized *sensitivity level* of the CSLN.

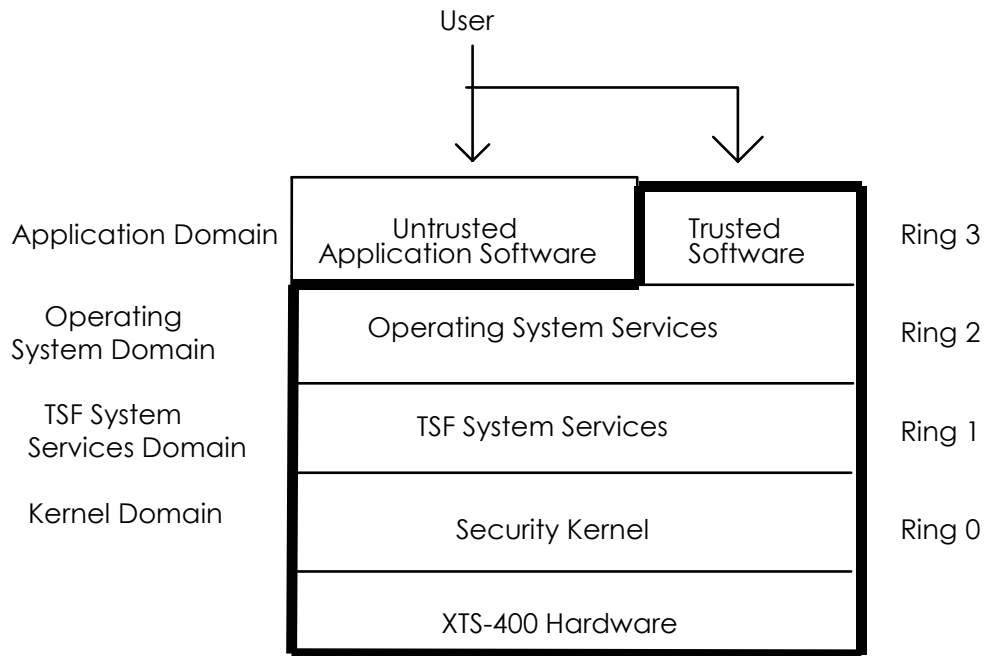


Figure 2. STOP System Diagram [7]

2.3 CSLN COMPONENT DESCRIPTIONS

The MYSEA CSLN is responsible for enabling protected simultaneous access between the MYSEA server and a large number of single-level networks. The MYSEA server, TCM and their ability to establish a *trusted channel* between themselves provides a distributed TSF perimeter and enables the MYSEA server to operate an MLS network interface. This architecture relies upon three principle components. The MYSEA server is the key component for this architecture and maintains primary responsibility for enforcing all security policies. The TCM is the second component and serves as a trusted endpoint for all *trusted channel* communications. An untrusted networking device is the third component and is required to multiplex a large number of CSLNs into one MYSEA server MLS network interface. Each single-level network may be considered an external component to the CSLN architecture. They will maintain the capacity to both send and receive data from the MYSEA server through the *trusted channel* via the TCM. The three components in this architecture are depicted in Figure 3.

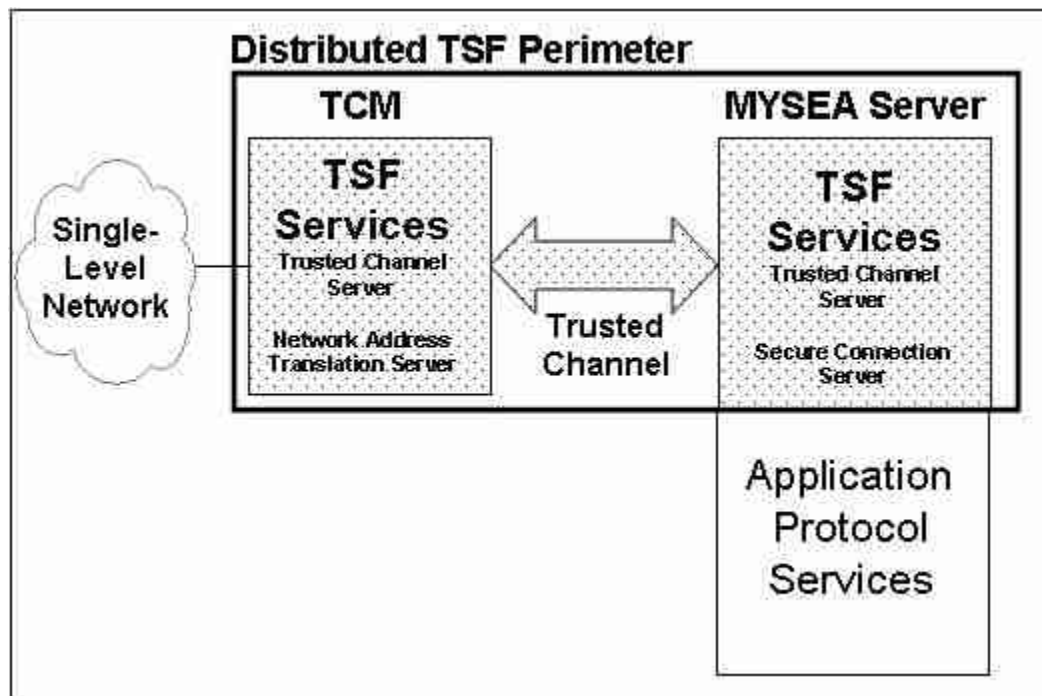


Figure 3. MYSEA CSLN Component Overview

2.3.1 MYSEA Server Target of Evaluation Security Function

The TSF “is a set consisting of all hardware, software, and firmware of the TOE that must be relied upon for the correct enforcement of the TSP” [2]. Although both the TCM and MYSEA server are designed with an individual TSF, this document will provide a design to create a distributed TSF for the MYSEA server that includes the TCM by establishing an *inter-TSF trusted channel* between the two security devices.

To create the security functions necessary to implement a distributed TSF, the MYSEA server requires: a Trusted Channel Server (TCS) to negotiate the security association of each PCC, authorize inbound communications and perform equivalence checks on the *sensitivity level* of all outbound communications; a Secure Connection Server (SCS) to ensure processes running on the server and their communications with the CSLN are authorized at the correct *sensitivity level*; a Trusted Channel Database (TCDB) to provide correlation between a TCM and the *sensitivity level* of the CSLN for which the TCM provides security services; a Secure Connection Inbound Database (SCIDB) to record the data required for the SCS to make its security determinations regarding each inbound connection; and a Secure Connection Outbound Database

(SCODB) to record the data required for the TCS to make its security determinations regarding each outbound connection. These trusted daemons and databases extend the TSF of the MYSEA server to the TCMs supporting each CSLN.

2.3.1.1 *Trusted Channel Server*

The TCS is a trusted daemon running in the *application domain* of the STOP. The TCS will perform two primary duties: 1) the TCS will negotiate the security association for each *trusted channel*; and 2) the TCS will authorize each inbound and outbound communication. The TCS will utilize a proprietary MYSEA Security Association protocol that is used for initializing, managing and terminating all *trusted channel* communications via the PCC protocol. The Security Association protocol will be similar to the Internet Key Exchange protocol [8], but will streamline the security association process to reduce the overall complexity of the protocol. The Security Association protocol will utilize digital certificates and predefined cryptographic algorithms to provide a secure means to negotiate the security parameters required to use the PCC protocol. High-level design specifications for the Security Association protocol are found in Appendix B of this thesis [9].

The TCS will also authorize each inbound and outbound communication. The TCS will work in concert with the PCC protocol handler to receive the implicit *sensitivity level* of each inbound connection so that it may query the TCDB for the explicit *sensitivity level* of each connection. The TCS will also check the SCODB for the explicit *sensitivity level* of each outbound connection. The TCS will use this data to query the TCDB for the *sensitivity level* of the requested single-level network and ensure that both *sensitivity levels* are equivalent before authorizing the connection.

Detailed high level design specifications for the TCS are found in Appendix C of this thesis [10].

2.3.1.2 *Secure Connection Server*

The SCS is a trusted daemon running in the STOP *Application Domain*. This process listens for TCS-authorized inbound CSLN connections. The SCS is responsible for checking the SCIDB for the explicit *sensitivity level* of each accepted connection and spawning the requested application protocol server at the corresponding *sensitivity level*. Additionally, the SCS will record the explicit *sensitivity level* of each

outbound connection to the SCODB. This data will be used by the TCS to validate that the *sensitivity level* of the requesting process and the *sensitivity level* of the requested single-level network are equivalent. The correct functionality of the SCS is imperative to enforce the TSP of the MYSEA server. SCS development parallels the “Secure Session Server” [1] with minor modifications to the original requirements and will be further defined in Appendix C of this thesis [10].

2.3.1.3 Trusted Channel Database

The TCDB is a static database used to maintain a list of all security attributes of all TCMs that are permitted to connect with the MYSEA server. Each entity in the database contain a record consisting of the Internet Protocol Routing Address (IP address) of each TCM, a listing of the subnets to which the TCM provides security services and its associated explicit *sensitivity level*. Only the TCS is permitted “read only” access to the TCDB. Trusted write operations to the TCDB will only be permitted by the security administrator in accordance with the Δ *property* first noted by Fellows [11]. Detailed high level design specifications for the TCDB are found in Appendix C of this thesis [10].

2.3.1.4 Secure Connection Inbound Database

The SCIDB is a dynamic database accessible only to the TCS and SCS. The TCS is responsible for creating a record for each new inbound connection and posting the explicit *sensitivity level* of each inbound connection to the newly created record. This record in-turn will be utilized by the SCS to spawn an application protocol server at the explicit *sensitivity level* of the inbound connection. The database permits read and write access for the TCS while permitting “read-only” access for the SCS. The TCS requires write access to the database so that a new record can be created for each new connection and so that the record can be deleted upon connection teardown. The TCS requires read access to the database so that a check can be accomplished to determine whether a record already exists for a new packet. The SCS requires read-only access to the database to query and return the explicit *sensitivity level* of each connection. The database maintains a record consisting of the original source IP address of the connection and the designated *sensitivity level* of the TCM. Detailed high level design specifications for the SCIDB are found in Appendix C of this thesis [10].

2.3.1.5 *Secure Connection Outbound Database*

The SCODB is a dynamic database accessible only to the SCS and TCS. The SCS is responsible for creating a new record in the SCODB consisting of the explicit *sensitivity level* of each outbound connection and the destination information. This database will provide the TCS with a basis for making security decisions on all outbound CSLN connections. The database permits read and *write* access for the SCS and the TCS. The SCS requires write access to the database so that a new record can be created to track each new connection. The SCS requires read access to the database so that a check can be accomplished to determine whether a record already exists for that packet. The TCS requires read access to the database so that it can obtain the explicit *sensitivity level* of each connection and requires write access to the database so that the record for each connection can be deleted upon connection teardown. The database maintains a record consisting of the destination IP address and the *sensitivity level* of the originating process. Detailed high level design specifications for the SCODB are found in Appendix C of this thesis [10].

2.3.1.6 *Protected Communications Channel Protocol Handler*

The PCC protocol is instrumental for establishing each *trusted channel*. The PCC concept was first introduced by Wilson [1] as a mechanism to establish a *trusted path* between the TPE and MYSEA server. The CSLN modifies and leverages the original PCC protocol to support a peer-to-peer *trusted channel* connection between the MYSEA server and the TCM. The PCC provides data segregation for all network traffic entering and exiting the MLS network interface. The PCC protocol handler provides *trusted channel* security and data segregation by providing the following security services: identification and authentication, integrity, and confidentiality.

Although these security services provide the necessary mechanisms to enable a *trusted channel*, they are not without residual risk. Basic traffic analysis within the CSLN will still be possible. Additionally, covert channels may possibly be created from the classified CSLNs. *Future work* on the MYSEA architecture should analyze these vulnerabilities and design changes to mitigate their risk. Appendix B of this thesis [9] provides detailed developmental analysis for the PCC protocol.

2.3.2 Trusted Channel Module Target of Evaluation Security Function Services

The TCM is a dedicated device placed on the network between each single-level network and the MYSEA server. The TCM will be engineered in accordance with Common Criteria security functions required for initiating *trusted channel* communications. The target Common Criteria Evaluation Assurance Level for the TCM will be EAL 6. To facilitate this goal, the TCM will leverage the Center for Information Systems Security Studies and Research (CISR) Trusted Computing Exemplar (TCX) Project [12] to provide a high assurance security kernel as its underlying operating system. The TCX kernel is intended to be evaluated at the EAL 7. The TCM must provide an unforgeable link between the MYSEA server and the CSLN. The TCM functionally serves as a high assurance endpoint virtual private network (VPN) device. The TCM will also serve as the front-end interface responsible for translating the IP addresses from its single level network domains to the MLS domain. The Department of Defense has allocated separate IP spaces for each of its unclassified and classified domains out of the worldwide IP address block. By policy, IP addresses in a higher classified domain can not be advertised and resolved in a lower classified domain [13]. To overcome IP address and resolution challenges for the MYSEA server, the TCM will serve as a network address translator for its CSLN. Information flows between the MYSEA server and the TCM have an implicit *sensitivity level*, which the MYSEA server utilizes to explicitly enforce its mandatory security policy. The TCM requires the development of two trusted daemons in order for it to provide a trusted endpoint with which the MYSEA server that can establish a *trusted channel*: The TCS and the Network Address Translation (NAT) [14] Server.

2.3.2.1 Trusted Channel Server

The TCS running on the TCM will mirror the functionality provided by the TCS running on the MYSEA server for initiating and managing the MYSEA Security Association protocol. The correct operation of the Security Association protocol is instrumental to the correct operation of the PCC protocol.

2.3.2.2 *Network Address Translation Server*

The NAT server will mediate all connectivity between the IP address space of the single-level network and the MYSEA IP address space. Rationale for providing NAT functionality for the CSLN can be found in the Chapter II of this thesis [15]. The NAT server will provide dynamic destination NAT on all inbound packets and static source NAT on all outbound packets.

2.3.2.3 *Protected Communications Channel Protocol Handler*

The PCC protocol handler will be embedded into the IP stack of the TCM and will be responsible for making all access control decisions from the single-level network to the MYSEA server and from the MYSEA server to the single-level network based upon its security policy database. The TCM implementation of the PCC protocol will be IPsec conformant.

2.3.3 Single Level Networks

Each single-level network represents a significant investment of resources (e.g. SIPRNET, JWICS and other compartmented or coalition partner networks) that will continue operations for the foreseeable future. MYSEA servers through their CSLN architecture will serve as a bridge between the various single-level networks until a comprehensive MLS architecture is fielded. Until then, single-level networks will continue to provide the backbone for MYSEA network connectivity, data and single *sensitivity level* applications.

3. SYSTEM REQUIREMENTS

3.1 CONNECTED SINGLE LEVEL NETWORK REQUIREMENTS

3.1.1 MYSEA shall be capable of supporting a large number of simultaneous single level network connections to and from multiple MYSEA servers.

3.1.2 MYSEA shall support these connections at multiple *sensitivity levels* with high assurance.

3.1.3 MYSEA shall provide identification and authentication, integrity and confidentiality to shared resources and application protocol services at designated *sensitivity levels* to each CSLN.

3.2 MYSEA SERVER REQUIREMENTS

The following MYSEA server requirements are necessary to manage simultaneous single level connections between itself and a TCM to establish a *trusted channel*. An abstract overview of a CSLN connection is presented in Figure 4.

3.2.1 The MYSEA server shall be able to establish multiple simultaneous *trusted channel* communications with pre-determined TCMs upon demand through a single MLS network interface.

3.2.2 The MYSEA server shall be able to support multiple MLS network interfaces.

3.2.3 Once a *trusted channel* is established with the TCM, any breakdown of the *trusted channel* detected by the MYSEA server shall lead to the termination of the connection independent of any other active *trusted channel* connections.

3.2.4 The MYSEA server shall associate a *sensitivity level* to a *trusted channel* based on the security attributes found in the PCC and further defined in a pre-configured trusted database.

3.2.5 The MYSEA server shall protect against disclosure and modification of information transiting the CSLN *trusted channels*.

3.2.6 The MYSEA server shall regulate all TCM access to itself.

3.2.7 The MYSEA server shall be responsible for *trusted channel* teardown upon the completion of each connection.

3.2.8 The MYSEA server shall implement network security mechanisms that protect against disclosure and modification of information transiting the *trusted channel*.

3.3 TRUSTED CHANNEL MODULE REQUIREMENTS

3.3.1 The TCM shall be able to establish multiple simultaneous *trusted channels*, all at the same *sensitivity level*, with multiple MYSEA servers.

3.3.2 Once a *trusted channel* is established, any breakdown of the *trusted channel* detected by the TCM shall lead to the termination of the connection.

3.3.3 The TCM shall have no runtime user interface.

3.3.4 The TCM shall implement network security mechanisms that protect against disclosure and modification of information transiting the *trusted channel*.

3.3.5 The TCM shall implement Network Address Translation services for all communications transiting the CSLN.

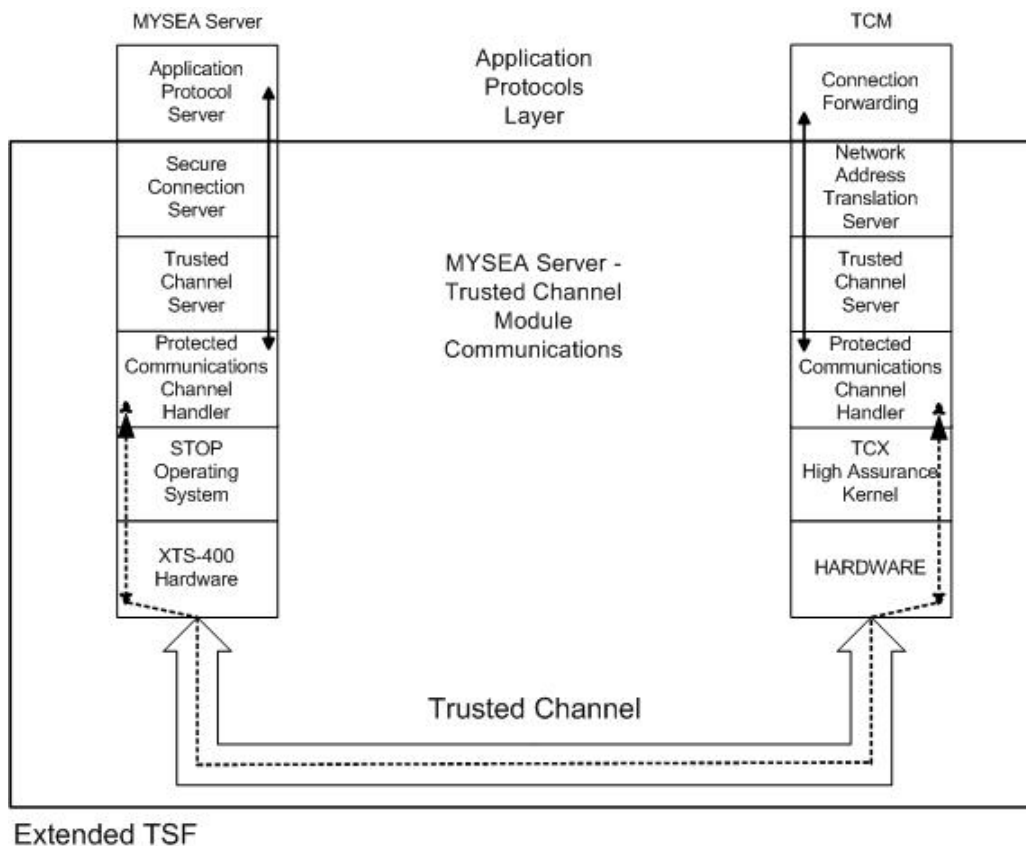


Figure 4. CSLN Connection Overview

3.4 MYSEA CONNECTION PROTOCOL REQUIREMENTS

3.4.1 MYSEA shall provide a communications protocol that facilitates a *logically distinct trusted channel* for the identification and authentication, integrity, and confidentiality of all authenticated communications between the MYSEA server and its authorized TCMs. Detailed requirements for the PCC protocol are found in Appendix B of this thesis [9].

3.5 CONNECTED SINGLE LEVEL NETWORK APPLICATION PROTOCOL SERVICES REQUIREMENTS

The CSLN network application protocol services requirements are enumerated below.

3.5.1 The CSLN shall have the capability to support all ISO Layer 3 and above application protocols.

3.5.2 The MYSEA server application protocol servers shall provide the capability to support commercial off the shelf and government off the shelf application products for authenticated users.

3.5.3 Access to resources and services on the MYSEA server from single level networks shall be controlled by the MYSEA server in accordance with its TSP.

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIXES

APPENDIX A ABBREVIATIONS, ACRONYMS, AND DEFINITIONS

A.1 Abbreviations, Acronyms

| | |
|---------|---|
| CISR | Center for Information Systems Security Studies and Research |
| CSLN | MYSEA Connected Single Level Network(s) |
| IP | Internet Protocol |
| JWICS | Joint Worldwide Intelligence Communications System |
| LAN | Local Area Network |
| MLS | Multilevel Security |
| MYSEA | Monterey Security Architecture |
| NIPRNET | NonSecure Internet Protocol Router Network |
| PCC | Protected Communications Channel |
| SCIDB | Secure Connection Inbound Database |
| SCODB | Secure Connection Outbound Database |
| SCS | Secure Connection Server |
| SIPRNET | SECRET Internet Protocol Router Network |
| STOP | Secure Trusted Operating Program |
| TCM | Trusted Channel Module |
| TCS | Trusted Channel Server |
| TCDB | Trusted Channel Server Database |
| TCX | Trusted Computing Exemplar |
| TOE | Target of Evaluation |
| TPE | Trusted Path Extension Device |
| TSF | Target of Evaluation Security Function |
| TSP | TOE Security Policy |

A.2 Definitions

2.1 Connected Single Level Network (CSLN): The segment of the MYSEA architecture from the MYSEA server to its supported TCMs responsible for providing single level network connectivity to the MYSEA server.

2.2 Inter-TSF Transfers: Communicating data between the TOE and the security functions of other trusted IT products [2].

2.3 Inter-TSF Trusted Channel: Requires that the TSF provide a trusted communication channel between itself and another trusted IT product [2].

2.4 Principle of Least Privilege: Every program and every user of the system should operate using the least set of privileges necessary to complete the job [6].

2.5 Sensitivity Level: The combined classification of data based upon its security or classification level and integrity level.

2.6 Target of Evaluation (TOE): An IT product or system and its associated guidance documentation that is the subject of an evaluation [2].

2.7 TOE Security Functions (TSF): A set consisting of all hardware, software, and firmware of the TOE that must be relied upon for the correct enforcement of the TSP [2].

2.8 TOE Security Functions Interface (TSFI): A set of interfaces, whether interactive (man-machine interface) or programmatic (application programming interface), through which TOE resources are accessed, mediated by the TSF, or information is obtained from the TSF [2].

2.9 TOE Security Policy (TSP): A set of rules that regulate how assets are managed, protected and distributed within a TOE [2].

2.10 Trusted Channel: A means by which a TSF and a remote trusted IT product can communicate with necessary confidence to support the TSP [2].

2.11 Trusted Channel Module: Security device required to enable “Inter-TSF Trusted Channels” between the MYSEA server and its authorized single level network.

2.12 Trusted Path: A means by which a user and a TSF can communicate with necessary confidence to support the TSP [2].

APPENDIX B. REFERENCES

- [1] J. D. Wilson, "A Trusted Connection Framework for Multilevel Secure Local Area Networks", Naval Postgraduate School, June 2000.
- [2] "Common Criteria for Information Technology Security Evaluation Volume I", *CCIMB-2004-01-001*, Version 2.2 ed: International Organization for Standardization, January 2004.
- [3] DigitalNet, "XTS-400 Trusted Computer System Technical Overview", http://www.digitalnet.com/solutions/information_assurance/pdf/XTS400%20Technical%20Description%206-8-04.PDF, Herndon, VA, July, 2004.
- [4] D. E. Bell and L. LaPadula, "Secure Computer System: Unified Exposition and Multics Interpretation", Mitre Corp., Hanscomb AFB, MA, Tech Rep ESD-TR-76-372, 1975.
- [5] K. J. Biba, "Integrity Considerations for Secure Computer Systems", MITRE Corp., Tech. Rep. ESD-TR-76-372, 1977.
- [6] J. H. Saltzer and M. D. Schroeder, "The Protection of Information in Computer Systems", *Proceedings of the IEEE*, Vol. 63. No. 9, Sept 1975, pp. 1278-1308.
- [7] DigitalNet, "XTS-400 User Manual STOP 6.0 Beta 12 Version", http://www.digitalnet.com/solutions/information_assurance/xts400_trusted_sys.htm, Herndon, VA, January 2003.
- [8] C. Kaufman, "The Internet Key Exchange (IKEv2) Protocol, draft-ietf-ipsec-ikev2-11.txt", *IP Security Protocol Working Group, IETF*, October 2003.
- [9] J. D. Sears, "Protected Communications Channel Protocol Requirements Document", Appendix B, "Simultaneous Connection Management and Protection in a Distributed Multilevel Security Environment", Masters Thesis, Naval Postgraduate School, Monterey, California, September 2004.
- [10] J. D. Sears, "Connected Single Level Network Trusted Channel Management Requirements Document", Appendix C, "Simultaneous Connection Management and Protection in a Distributed Multilevel Security Environment", Masters Thesis, Naval Postgraduate School, Monterey, California, September 2004.
- [11] J. Fellows, J. Hemenway, N. Kelem, and S. Romero, "The Architecture of a Distributed Trusted Computing Base", *Proceedings of the 10th National Conference on Computer Security*, pp. 68-77, September 1987.

- [12] Irvine, Cynthia E., Levin, Timothy E., Nguyen, Thuy D., and Dinolt, George W., "The Trusted Computing Exemplar Project", Proceedings of the 2004 IEEE Systems Man and Cybernetics Information Assurance Workshop, West Point, NY, June 2004, pp. 109 - 115.
- [13] Defense Information Systems Agency, "SIPRNET Classification Guide."
- [14] K. Srisuresh and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT), RFC 3022," *Network Working Group, IETF*, January 2001.
- [15] J. D. Sears, "Simultaneous Connection Management and Protection in a Distributed Multilevel Security Environment", Masters Thesis, Naval Postgraduate School, Monterey, California, September 2004.

APPENDIX B

PROTECTED COMMUNICATIONS CHANNEL PROTOCOL REQUIREMENTS DOCUMENT

TABLE OF CONTENTS

| | | |
|-----|--|----|
| 1. | INTRODUCTION..... | 69 |
| 1.1 | PURPOSE..... | 69 |
| 1.2 | SCOPE | 69 |
| 2. | PROTECTED COMMUNICATIONS CHANNEL OVERVIEW | 71 |
| 2.1 | MYSEA CONNECTIVITY..... | 71 |
| 2.2 | CSLN TRANSMISSION SECURITY | 71 |
| 2.3 | TRUSTED CHANNEL COMMUNICATIONS..... | 71 |
| 2.4 | PROTECTED COMMUNICATIONS CHANNEL PROTOCOL..... | 73 |
| 2.5 | MYSEA SERVER APPLICATION PROTOCOL SERVICES | 74 |
| 3. | CONNECTION PROTOCOL REQUIREMENTS | 77 |
| 3.1 | PROTECTED COMMUNICATIONS PROTOCOL REQUIREMENTS..... | 77 |
| 3.2 | IPSEC PROTOCOL REQUIREMENTS | 77 |
| | APPENDIXES | 83 |
| | APPENDIX A ABBREVIATIONS, ACRONYMS, AND DEFINITIONS..... | 83 |
| | APPENDIX B REFERENCES | 85 |

THIS PAGE INTENTIONALLY LEFT BLANK

1. INTRODUCTION

1.1 PURPOSE

The purpose of this systems requirements document is to provide a high level analysis of the incorporation of IPsec into the Monterey Security Architecture (MYSEA) Protected Communications Channel (PCC) Protocol and describe the IPsec functionality necessary to enable a *trusted channel*. Requirements for the PCC were initially identified in “A Trusted Connection Framework for Multilevel Secure Local Area Networks” [1]. Common Criteria standards and terminology shall be used to describe all hardware, software, firmware, and their interactions [2].

1.2 SCOPE

This document defines the MYSEA PCC protocol functionality required to establish a *trusted channel* between the MYSEA server and a new security device hereafter known as the Trusted Channel Module (TCM). The PCC protocol also applies to a *trusted path* between the MYSEA server and the Trusted Path Extension (TPE) device. The integration of a *trusted channel* into MYSEA is crucial in order for the MYSEA server to provide protected, simultaneous connections to a large number of single level networks (e.g., NIPRNET, SIPRNET, JWICS) through a single multilevel secure (MLS) network interface. These connections will operate at multiple *sensitivity levels* while providing simultaneous management of all connections. Satisfying these requirements provides MYSEA with a required component necessary to provide a robust distributed MLS environment.

THIS PAGE INTENTIONALLY LEFT BLANK

2. PROTECTED COMMUNICATIONS CHANNEL OVERVIEW

2.1 MYSEA CONNECTIVITY

Producing a distributed (MLS environment featuring commercial-off-the-shelf form and features is the objective of MYSEA. This architecture features two distinct segments that enable MLS connectivity to the end user. The MYSEA Local Area Network (LAN) connects the end user to the MYSEA MLS server and is capable of delivering multiple *sensitivity levels* of data and applications to the authorized end user. The Connected Single Level Network (CSLN) affords the MYSEA server a mechanism for protected connectivity to multiple simultaneous single level networks by providing a distributed Trusted Security Function (TSF) to the TCM. Currently, the MYSEA server has a limited number of network interfaces to connect this large number of single level networks. The Department of Defense and private industry support a very large number of system high, dedicated and compartment single level networks, each operating with multiple compartments, caveats, and releasability issues. Therefore, evolving the functionality to include a MLS network interface into the MYSEA server becomes an essential feature for providing a truly distributed MLS environment. Figure 1 presents an overview of CSLN communications.

2.2 CSLN TRANSMISSION SECURITY

Transmission security (TRANSEC) of all *trusted channel* communications between the MYSEA server and the TCM is the critical factor required to establish a distributed MLS architecture. An unforgeable link between these two critical-security components protecting the integrity and confidentiality of all communications while providing absolute identification and authentication is an unconditional requirement for establishing this *trusted channel*.

2.3 TRUSTED CHANNEL COMMUNICATIONS

Establishing a *trusted channel* is the key enabling operation for CSLN communications. The *trusted channel* enables the MYSEA server to ensure only authorized communications are permitted into the MYSEA server from the TCMs and from the MYSEA server through the TCMs to the protected enclave of each single-level




Figure 1. MYSEA Connected Single Level Network Overview

network. A *trusted channel* is defined in Section 13 of The Common Criteria [3] as “a trusted communications channel between the TSF and other trusted IT products.” It further states that “a *trusted channel* is a communication channel that may be initiated by either side of the channel, and provides non-repudiation characteristics with respect to the identity of the sides of the channel.” FTP_ITC Inter-TSF Trusted Channels is the governing Common Criteria subsection mandating required *trusted channel* behaviors. The following behaviors are defined for MYSEA CSLNs:

2.3.1 FTP_ITC.1.1 - The TSF shall provide a communication channel between itself and a remote trusted IT product that is logically distinct from other communication channels and provides assured identification of its end points and protection of the channel data from modification or disclosure.

2.3.2 FTP_ITC.1.2 - The TSF shall permit both the TSF and the remote trusted IT product (i.e., TCM) to initiate communication via the *trusted channel*.

2.3.3 ¹ FTP_ITC.1.3 - The TSF shall use a *trusted channel* for all communications between the MYSEA server and TCM.

¹ FTP_ITC 1.3 Security Functional Requirement modified to include Common Criteria Observation Decisions Review Board Precedent Database revision to eliminate contradiction between FTP ITC 1.2 and FTP ITC 1.3.

2.4 PROTECTED COMMUNICATIONS CHANNEL PROTOCOL

The CSLN encompasses all *trusted channel* initialization and *trusted channel* communications between the MYSEA server and its authorized TCMs. The CSLN will utilize a cryptographic communications protocol to establish high assurance, logically distinct communications channels providing authentication, integrity and confidentiality. The CSLN will build upon earlier PCC protocol analysis [1] and further define its required attributes to invoke a *trusted channel*. The PCC will be integrated into the IP stack of the host operating system. A trusted daemon hereafter known as the Trusted Channel Server (TCS) will be responsible for initiating, managing, and terminating each PCC connection. The TCS will operate on both the MYSEA server and the TCM.

The PCC packet carries with it an implied *sensitivity level* for its CSLN. The PCC protocol handler will implement a custom API to extract the source IP addresses from the inner and outer IPsec headers and pass that data to the TCS for further security-related actions. The TCS must acknowledge that the transfer was successful before the PCC process can continue. The TCS works in conjunction with the Secure Connection Server (SCS) and three protected databases known as the Trusted Channel Database (TCDB), the Secure Connection Outbound Database (SCODB) and the Secure Connection Inbound Database (SCIDB) to multiplex a very large number of CSLNs into one MYSEA server MLS network interface. High level design requirements for these daemons and databases are found in Appendix A. Detailed implementation design for these daemons and databases are found in Appendix C.

Although the PCC will be encrypted, initial iterations of the protocol will not meet Type I requirements as mandated by the Department of Defense [4]. Therefore, National Security Agency Type I certified encryption equipment shall be used to cover any classified communications between the TCM and MYSEA server where the TCM is not physically co-located with the MYSEA server in a protected environment.

2.4.1 IPsec Protocol – MYSEA selected IPsec as its PCC protocol after analyzing the strengths and weaknesses of the various candidate protocols [1]. IPsec supports multiple transport protocols, cryptographic algorithms and key management schemes, which yields the necessary flexibility to ensure the PCC delivers a full range of security services to MYSEA. In accordance with Common Criteria requirements

specified for implementing *inter-TSF trusted channels*, IPsec supports initiation of the *trusted channel* by either side of the channel, identification and authentication, integrity, and confidentiality for each transmitted packet. Additionally, IPsec provides non-repudiation, logically distinct data separation and anti-replay protection. IPsec enables implementation of access control lists (ACL), which provides the fine grain control necessary to authorize all inbound and outbound communications with router-equivalent specificity. Requirements for the MYSEA IPsec implementation are detailed in Section 3.2 of this appendix and are founded upon Draft IETF IPsec RFC2401bis-01, “Security Architecture for the Internet Protocol” [5].

2.4.2 Residual Risk – The security services provided by IPsec afford the CSLN architecture the mechanism needed to establish a *trusted channel* between the MYSEA server and TCM. However, some residual risk remains. First, IPsec is limited to using Type II cryptographic algorithms, which are not certified for use with classified data. Second, IPsec security mechanisms cannot stop basic traffic analysis techniques. Even so, the DoD has already accepted this risk through the use of its current generation of Type I IP encryptors. Third, the possibility exists for covert channels to be created in the CSLN. *Future work* for MYSEA should include a residual risk study of the CSLN and design changes to mitigate those risks.

2.5 MYSEA SERVER APPLICATION PROTOCOL SERVICES

Once the TCM and MYSEA server are connected via the *trusted channel*, authorization to access application-level protocol services (e.g., IMAP, FTP, HTTP, etc.) may be granted to the single-level network based upon the sensitivity level of the connection. The TCS shall be responsible for querying the TCDB and performing this security-critical operation. Upon authentication, the TCS shall post the *sensitivity level* of each connection to the SCIDB. The SCS shall query the SCIDB for the *sensitivity level* of the connections and spawn the requested application protocol server [6] at the correct *sensitivity level* as determined by the TCS. Note that a connection from a single-level network to the MYSEA server is only authenticated to the network level, not the individual user. *Future work* for the MYSEA architecture should include mechanisms for authenticating individual single-level network users to the MYSEA server.

Application service validation mechanisms will be covered in detail in Appendix C of this document [7].

THIS PAGE INTENTIONALLY LEFT BLANK

3. CONNECTION PROTOCOL REQUIREMENTS

3.1 PROTECTED COMMUNICATIONS PROTOCOL REQUIREMENTS

3.1.1 Both the MYSEA server and the TCM shall have the capability to initiate the establishment of the *trusted channel*.

3.1.2 The *trusted channel* shall provide secure communications to include authentication, integrity and confidentiality between the MYSEA server and the TCM.

3.1.3 The *trusted channel* shall implement anti-replay control features.

3.1.4 The *trusted channel* shall utilize the IPsec protocol as the underlying mechanism to implement the PCC protocol.

3.1.5 The TCS shall initiate and manage all PCC connections.

3.1.6 Explicit end of connection tear-down commands shall be sent only by the MYSEA server.

3.1.7 Upon recognition of a connection failure by either the MYSEA server or the TCM, the recognizing component shall tear-down the connection – *fail secure* [8].

3.1.7.1 Only authenticated communications shall transit the CSLN between the MYSEA server and its associated TCMs.

3.2 IPSEC PROTOCOL REQUIREMENTS

3.2.1 Encapsulating Security Payload (ESP) protocol - ESP shall be invoked to provide packet authentication and confidentiality. Although projected implementations of the CSLN will connect the MYSEA server to a TCM co-located in a secure environment, ESP remains a requirement due to the mandatory inclusion of an untrusted networking device between the MYSEA server and its TCM. This device is responsible for multiplexing a large number of TCM connections into one MYSEA server network interface. To prevent the possible exploitation of MLS data transiting this low assurance, untrusted device, ESP encryption will be used to enforce data separation and confidentiality of all trusted communications while in transit between the MYSEA server and TCM. ESP protects against IP spoofing, packet modification and unauthorized disclosure between the MLS server and the CSLN client.

3.2.2 Authentication Header (AH) protocol - AH shall be invoked to provide packet authentication and integrity to include the immutable IP header fields of the outer header and anti-replay protection. Although many of the security functions between ESP and AH appear the same, the ability of AH to provide data integrity on the outer header is a crucial design feature necessary to invoke a *trusted channel*. This requirement will be discussed in detail in Appendix C. AH protects against IP spoofing, packet modification, and replay attacks between the MLS server and the TCM.

3.2.3 IPsec Physical Implementation - Integrating IPsec may occur at three layers to include [5]:

3.2.3.1 Integration in the native IP stack. Native integration permits IPsec to be embedded directly into the IP layer source code. This method of integration provides for more efficient processing than the other two methods of IPsec integration, both of which will be discussed in the next section. However, *Native* integration may be the most difficult for the MYSEA project to implement, as it requires access to the source code of the operating system.

3.2.3.2 Integration as a "bump-in-the-stack" (BITS).

IPsec is implemented "underneath" an existing implementation of an IP protocol stack, between the native IP and the local network drivers. Source code access for the IP stack is not required in this context, making this implementation approach appropriate for use with legacy systems. This approach, when it is adopted, is usually employed in hosts.

3.2.3.3 Integration as a "bump-in-the-wire" (BITW).

The use of a dedicated, inline security protocol processor is a common design feature of systems used by the military, and of some commercial systems as well. It is sometimes referred to as a "bump-in-the-wire" (BITW) implementation. Such implementations may be designed to serve either a host or a gateway. Usually the BITW device is itself IP addressable. When supporting a single host, it may be quite analogous to a BITS implementation, but in supporting a router or firewall, it must operate like a security gateway.

Due to the *gateway* functionality of the TCM, the TCM shall implement a BITW IPsec configuration in accordance with the RFC definition. Integrating IPsec into the MYSEA server is more complicated. If possible, a *native* integration into the MYSEA server shall be used, which yields the best possible results. If integration at the *native* layer is not possible, a BITS IPsec implementation shall be used in the MYSEA server.

3.2.4 Security Association and Key Management - IPsec includes the capability to perform manual or automated security association (SA) and cryptographic key management. Internet Key Exchange Version 2 (IKE2) [9] is the preferred IPsec automated SA and cryptographic key management protocol. However, MYSEA shall implement a custom automated SA protocol that may inherit features from IKE2, but will be substantially reduced in complexity. This new protocol will streamline the connection negotiation process and minimize the complexity necessary to evaluate the correctness of the automated SA protocol for a high assurance evaluation. The MYSEA automated SA and key management architecture shall implement the following requirements.

3.2.4.1 Digital Certificates. MYSEA shall use digital certificates to assign a public/private key pair to the MYSEA server and its associated TCMs. All initialization messages shall be signed with the private key owned by the sender and encrypted with the public key owned by the receiver. Upon receipt of each message, the receiver shall validate that the request is from an authorized sender using the public key owned by the sender and decrypt the data using the private key owned by the receiver.

3.2.4.2 Key Negotiation and Management. The MYSEA server shall be solely responsible for generating a unique symmetric connection key for each *trusted channel*. Each connection key shall be sent to the TCM and used for communications specific to the negotiated connection.

3.2.4.3 IPsec Algorithms. IPsec includes several cryptographic algorithms to enforce its authentication, confidentiality and integrity mechanisms. MYSEA shall use two predetermined algorithms to eliminate the complexity necessary to negotiate the cryptographic algorithms of each connection. MYSEA shall use the AES-128-CBC [10] algorithm to enforce confidentiality and the HMAC-SHA1 [11] algorithm

to enforce integrity. Should these algorithms be found to have vulnerabilities, alternative algorithms may be substituted in their place.

3.2.4.4 Security Parameter Index (SPI). The MYSEA server and TCM shall each create a unique SPI upon initial connection negotiation so that anti-replay protection features are enabled.

A simplified overview of the TCM-to-MYSEA server automated SA process is illustrated in Figure 2. The TCM shall send a connection request to the MYSEA server. The MYSEA server shall then create an initial security parameter index (SPI) and symmetric connection key unique to the connection and send the data to the requesting TCM. The TCM shall acknowledge the connection negotiation data from the server, reply with its own SPI and establish a SA based upon the symmetric key created by the server. The server will acknowledge the SPI from the TCM, create a SA for the connection and declare that it is ready to receive. The TCM will commence data transmission with the negotiated symmetric connection key.

A simplified overview of the MYSEA server-to-TCM automated SA process is illustrated in Figure 3. The MYSEA server shall create an initial security parameter index and symmetric connection key unique to the connection and send the data to the requested TCM. The TCM shall acknowledge the connection negotiation data from the server, create and reply with its own SPI and establish a SA based upon the symmetric key created by the server. The server will acknowledge the SPI from the TCM and create a SA for the connection. The MYSEA server will commence data transmission with the negotiated symmetric connection key.

Specific design and development details regarding the MYSEA automated SA protocol is left for *future work*.

3.2.5 Management of Covert Channels - IPsec *tunnel mode* recognizes that mutable differentiated services code point fields (DSCP) may be used to provide covert communications. MYSEA shall map the DSCP field to a fixed value to negate this potential vulnerability.

3.2.6 MYSEA integration of the IPsec Protocol - Per Appendix A, the MYSEA TCS shall associate an explicit *sensitivity level* to all incoming *trusted channel* communications based upon the IP address of the authenticated TCM. Using IPsec

tunnel mode, the TCM IP address is contained in the outer IP header which is stripped away during initial IPsec inbound processing. This leaves only the original source IP address for processing at the network and transport layers. Therefore, the PCC protocol handler shall provide a custom API to extract and return the inner and outer source IP addresses located in the IPsec header and send that data to the TCS for further security decisions.

3.2.7 Future work - Initial iterations of the TCM shall provide no remote user interface. Future iterations may include an IPsec *transport* mode implementation to facilitate a *Simple Network Management Protocol* [12] equivalent for remote TCM administration and auditing review.

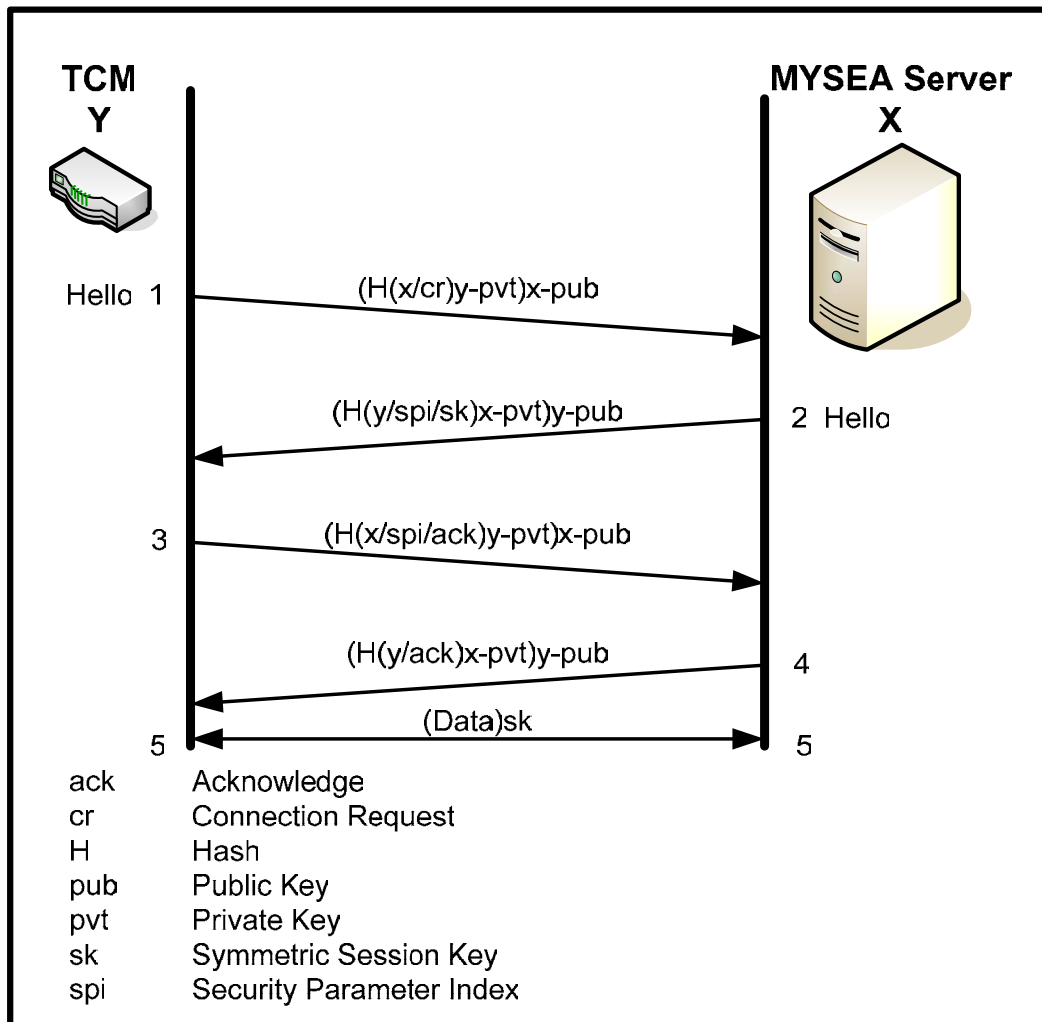


Figure 2. TCM-to-MYSEA Server Automated Security Association

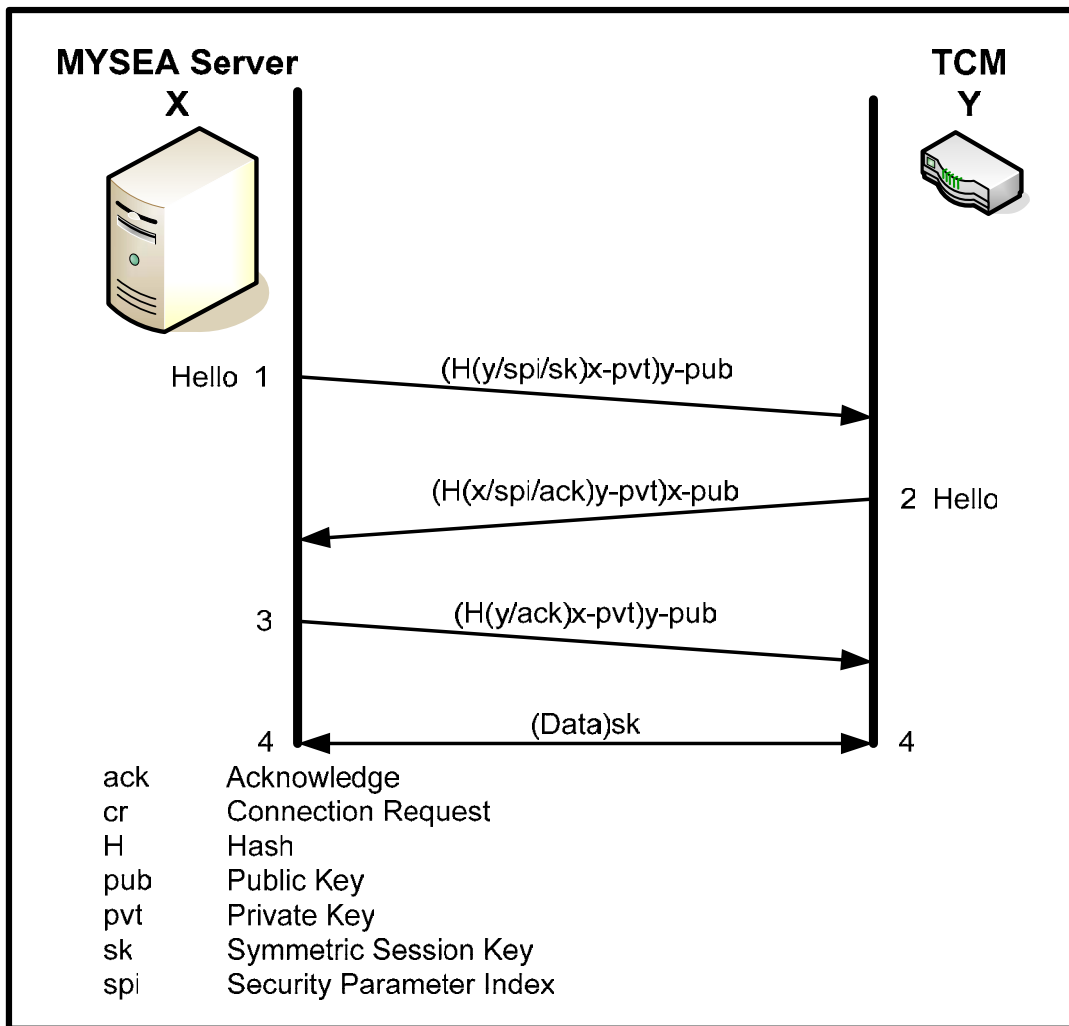


Figure 3. MYSEA-to-TCM Server Automated Security Association

APPENDIXES

APPENDIX A ABBREVIATIONS, ACRONYMS, AND DEFINITIONS

A.1 Abbreviations, Acronyms

| | |
|---------|---|
| ACL | Access Control List |
| AH | Authentication Header Protocol |
| BITS | Bump-in-the-Stack |
| BITW | Bump-in-the-Wire |
| CSLN | Connected Single Level Network |
| ESP | Encapsulating Security Payload Protocol |
| IKE2 | Internet Key Exchange Protocol Version 2 |
| IP | Internet Protocol |
| JWICS | Joint Worldwide Intelligence Communications System |
| LAN | Local Area Network |
| MLS | Multilevel Security |
| MYSEA | Monterey Security Architecture |
| NIPRNET | NonSecure Internet Protocol Router Network |
| PCC | Protected Communications Channel |
| RFC | Request for Comment |
| SA | Security Association |
| SCS | Secure Connection Server |
| SCIDB | Secure Connection Inbound Database |
| SIPRNET | SECRET Internet Protocol Router Network |
| SCODB | Secure Connection Outbound Database |
| SPI | Security Parameter Index |
| TCM | Trusted Channel Module |
| TCS | Trusted Channel Server |
| TCDB | Trusted Channel Database |

| | |
|---------|--|
| TRANSEC | Transmission Security |
| TSF | Target of Evaluation Security Function |

A.2. Definitions

2.1 Connected Single Level Network (CSLN): The segment of the MYSEA architecture from the MYSEA server to its supported TCMs responsible for providing single level network connectivity to the MYSEA server.

2.2 Fail-secure: Fail secure asserts that no compromise occurs even when some components are unavailable [8].

2.3 Sensitivity Level: The combined classification of data based upon its security or classification level and integrity level.

2.4 Trusted Channel: A means by which a TSF and a remote trusted IT product can communicate with necessary confidence to support the Target of Evaluation Security Policy [3]

2.5 Trusted Channel Module (TCM): Security device required to enable “Inter-TSF Trusted Channels” with the MYSEA server.

2.6 Target of Evaluation (TOE): An IT product or system and its associated guidance documentation that is the subject of an evaluation [2].

2.7 TOE Security Functions (TSF): A set consisting of all hardware, software, and firmware of the TOE that must be relied upon for the correct enforcement of the TSP [2].

2.8 Transmission Security (TRANSEC): The component of Communication Security (COMSEC) that results from the application of measures designed to protect transmissions from interception and exploitation by means other than cryptanalysis [4].

APPENDIX B REFERENCES

- [1] J. D. Wilson, "A Trusted Connection Framework for MultiLevel Secure Local Area Networks ", Naval Postgraduate School, June 2000.
- [2] "ISO/IEC 14508 - Common Criteria for Information Technology Security Evaluation Volume I," *CCIMB-2004-01-001*, Version 2.2 ed: International Organization for Standardization, January 2004.
- [3] "ISO/IEC 14508 - Common Criteria for Information Technology Security Evaluation Volume II," *CCIMB-2004-01-002*, Version 2.2 ed: International Organization for Standardization, January 2004.
- [4] Department of Defense, "Directive C-5200.5 Communications Security," 21 April 1990.
- [5] S. Kent and K. Seo, "Security Architecture for the Internet Protocol, Draft-IETF-IPsec-RFC2401BIS-01," *IP Security Protocol Working Group, IETF*, July 2004.
- [6] S. Bryer-Joyner and S. Heller, "Secure Local Area Network Services for a High Assurance Multilevel Network ", Naval Postgraduate School, Monterey, CA. March 1999.
- [7] J. D. Sears, "Connected Single Level Network Trusted Channel Management Requirements Document ", Naval Postgraduate School, September 2004.
- [8] J. Fellows, J. Hemenway, N. Kelem, and S. Romero, "The Architecture of a Distributed Trusted Computing Base," *Proceedings of the 10th National Conference on Computer Security*, pp. 68-77, September 1987.
- [9] C. Kaufman, "The Internet Key Exchange (IKEv2) Protocol, draft-ietf-ipsec-ikev2-11.txt," *IP Security Protocol Working Group, IETF*, October 2003.
- [10] National Institute of Standards and Technology, "Federal Information Processing Standards Publication 197 Announcing the Advanced Encryption Standard," <http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf>, November 2001.
- [11] National Institute of Standards and Technology, "Federal Information Processing Standards Publication 180-2 Announcing the Secure Hash Standard," <http://csrc.nist.gov/publications/fips/fips180-2/fips180-2.pdf>, August 2002.
- [12] J. Case, M. Fedor, M. Schoffstall, and J. Davin, "Simple Network Managment Protocol (SNMP)," *Network Working Group, IETF*, May 1990.

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX C

CONNECTED SINGLE LEVEL NETWORK TRUSTED CHANNEL MANAGEMENT REQUIREMENTS DOCUMENT

TABLE OF CONTENTS

| | | |
|-------|--|-----|
| 1. | INTRODUCTION..... | 89 |
| 1.1 | PURPOSE..... | 89 |
| 1.2 | SCOPE | 89 |
| 2. | TRUSTED CHANNEL MANAGEMENT | 91 |
| 2.1 | OVERVIEW | 91 |
| 2.2 | PROTECTED COMMUNICATIONS CHANNEL PROTOCOL..... | 92 |
| 2.3 | MYSEA SERVER..... | 93 |
| 2.3.1 | Trusted Channel Server | 93 |
| 2.3.2 | Secure Connection Server | 96 |
| 2.3.3 | Trusted Channel Database..... | 97 |
| 2.3.4 | Secure Connection Inbound Database..... | 98 |
| 2.3.5 | Secure Connection Outbound Database | 98 |
| 2.4 | TRUSTED CHANNEL MODULE..... | 98 |
| 2.4.1 | Trusted Channel Server | 98 |
| 2.4.2 | Network Address Translation Server | 98 |
| 3. | TRUSTED CHANNEL MODULE TO MYSEA SERVER MANAGEMENT FUNCTIONS..... | 99 |
| 3.1 | OVERVIEW..... | 99 |
| 3.2 | TCM INBOUND STATES | 99 |
| 3.3 | MYSEA SERVER CONNECTED SINGLE LEVEL NETWORK INBOUND STATES | 102 |
| 4. | MYSEA SERVER TO TRUSTED CHANNEL MODULE MANAGEMENT FUNCTIONS..... | 107 |
| 4.1 | OVERVIEW..... | 107 |
| 4.2 | MYSEA SERVER CONNECTED SINGLE LEVEL NETWORK OUTBOUND STATES | 107 |
| 4.3 | TCM OUTBOUND STATES..... | 111 |
| | APPENDIXES..... | 115 |
| | APPENDIX A ABBREVIATIONS, ACRONYMS, AND DEFINITIONS..... | 115 |
| | APPENDIX B REFERENCES | 117 |

THIS PAGE INTENTIONALLY LEFT BLANK

1. INTRODUCTION

1.1 PURPOSE

Providing a specific framework for enabling the Connected Single Level Network (CSLN) segment of the Monterey Security Architecture (MYSEA) is the purpose of this document. This framework extends the initial requirements for MYSEA first identified in Appendix C of “A Trusted Connection Framework for Multilevel Secure Local Area Networks” [1] and later expanded upon in an “Overview of High Assurance Architecture for Distributed Multilevel Security” [2]. Common Criteria standards and terminology shall be used to describe all hardware, software, firmware and their interactions [3].

1.2 SCOPE

MYSEA provides a framework for supporting multilevel security (MLS) to end user clients with commercial-off-the-shelf functionality. The MYSEA CSLN segment encompasses all communications between two high assurance devices. The first device is the MYSEA server which uses the DigitalNet XTS-400 Trusted Computing System running the Secure Trusted Operating Program (STOP) [4]. The second device is the Trusted Channel Module (TCM) under development at the Center for Information Systems Security Studies and Research (CISR) at the Naval Postgraduate School. These two high assurance devices provide the underlying security functionality required to create a *trusted channel*, which enables the MYSEA server to multiplex a large number of existing system high, compartmented and dedicated networks (e.g., NIPRNET, SIPRNET, JWICS) into a single MLS network interface. This document will provide detailed explanations concerning the *trusted channel* protocols, management functions and state transitions between the Trusted Channel Module (TCM) and the MYSEA server required to facilitate the CSLN segment.

THIS PAGE INTENTIONALLY LEFT BLANK

2. TRUSTED CHANNEL MANAGEMENT

2.1 OVERVIEW

The CSLN segment depends upon several modifications to the existing MYSEA architecture. Most notably, a *trusted channel* between the MYSEA server and TCM must be created to affect the required functionality necessary to enable MLS network interfaces. MYSEA designers designated the protocol required to implement a *trusted channel* the Protected Communications Channel (PCC) protocol. Additionally, the MYSEA server requires the creation of two trusted daemons, the Trusted Channel Server (TCS) and the Secure Connection Server (SCS). These daemons extrapolate an implied *sensitivity level* associated with a PCC connection and enforce an explicit *sensitivity level* on all *trusted channel* connections processed by the MYSEA server. The TCS is responsible for setting-up and terminating PCC connections as well as associating a *sensitivity level* with all inbound and outbound PCC connections. For incoming connections, the SCS is responsible for spawning application protocol servers at the correct *sensitivity level*. For outgoing connections, the SCS is responsible for generating the required information necessary for the TCS to enforce the overall security policy of the server.

The creation of three protected databases is also required in the MYSEA server: 1) the Trusted Channel Database (TCDB); 2) the Secure Connection Inbound Database (SCIDB); and 3) the Secure Connection Outbound Database (SCODB). These three databases work in concert with the TCS and SCS to track the *sensitivity level* of all *trusted channel* communications. The TCDB is a protected static database which associates an explicit *sensitivity level* between a TCM and its CSLN IP address. The SCIDB is a protected database used to record the explicit *sensitivity level* of each inbound connection from the *trusted channel*. The records in the SCIDB are used to ensure that, for each trusted channel, the sensitivity level of the server process to which the packet is destined is equal to that of the connection. The SCODB is a protected database used to verify that each MYSEA server application requesting connections with a CSLN has an equivalent *sensitivity level*. The records in the SCODB are used by the TCS to enforce this *sensitivity level* check.

The TCM requires the implementation of two trusted daemons to manage the *trusted channel*: The first is the TCM TCS; and the second is the TCM Network Address Translation (NAT) server. The TCM TCS will work in step with the MYSEA server to set-up and terminate each PCC connection. The NAT server will provide destination NAT on all CSLN inbound connections and source NAT on all outbound CSLN connections. The rationale for imposing NAT functionality in the TCM stems from the current DoD IP address security policy. The DoD owns a significant block of the global IP address spaces. Current Department of Defense policy stipulates the segmentation of this IP space between its various unclassified and classified domains to ensure that no overlap of the global IP address space exists. The Network Information Center at the Defense Information Systems Agency (DISA) is the primary entity responsible for allocating the DoD IP space. The DISA “SIPRNET Classification Guide” [5] stipulates that classified domain IP addresses remain physically and cryptographically separated from the unclassified IP domain and, by policy, a classified IP address can not be associated with a place or system, or advertised in a domain of lower classification. As a result, MYSEA servers acting in their MLS capacity will be required to operate in a separate IP space until a DoD MLS IP policy can be formulated.

2.2 PROTECTED COMMUNICATIONS CHANNEL PROTOCOL

Per research presented in “A Trusted Connection Framework for Multilevel Secure Local Area Networks” [1], the PCC shall be based on the IPsec protocol. High level design and implementation details are specified in Appendix B of this thesis [6]. The PCC shall provide authentication, confidentiality and integrity for all CSLN data transiting between the MYSEA server and the TCM. These security services will be enabled by evoking the IP Encapsulating Security Payload (ESP) protocol and the IP Authentication Header (AH) protocol in *Tunnel Mode*.

The ESP protocol will be used to protect the confidentiality of all CSLN connections and ensure data segregation between the MYSEA server and its associated TCMs. The AH protocol will be used to protect the authenticity and integrity of all CSLN connections. Although ESP has an integrity protection mechanism, that mechanism does not protect the outer IPsec header created when using tunnel mode. As such, ESP integrity protection is not used by the PCC protocol. Hence, AH is required to

provide integrity across the IPsec packet to include the immutable fields of the new outer IP header. *Tunnel Mode* is an IPsec mode of operation used to protect packets traversing a security gateway device like the TCM. Tunnel mode encapsulates the original IP packet in a new IPsec packet and incorporates a new IP header and several additional fields necessary to complete each IPsec transaction. ESP and AH each have unique fields that are applied as each protocol is invoked. These protocols can be applied sequentially on the same packet, known as a *security bundle*, to layer both ESP and AH security services on the packet which they are acting upon.

Although the IPsec protocol provides authentication, confidentiality and integrity, some residual risk remains. IPsec cannot mitigate the vulnerabilities associated with basic traffic analysis and covert channels. Future work on MYSEA should include a thorough risk assessment of the CSLN architecture with respect to these vulnerabilities.

2.3 MYSEA SERVER

In this section, the MYSEA server trusted channel management functions necessary to create an MLS network interface are presented.

2.3.1 Trusted Channel Server

The TCS will be a trusted multi-function daemon instrumental to the enforcement of the Target of Evaluation Security Policy (TSP) of the MYSEA server. The TCS will be designed to provide two critical security services: 1) the TCS will manage a proprietary MYSEA Security Association protocol which is essential to the secure setup of the PCC (IPsec) connection; and 2) the TCS will be responsible for enforcing the TSP of the MYSEA server by authorizing each inbound and outbound connection based upon its explicit *sensitivity level*.

The TCS will set-up and terminate all PCC connections. Upon recognition of a new PCC connection request, the PCC protocol handler will call the TCS to set-up the connection by applying the MYSEA Security Association protocol. The high level design of this protocol is found in Appendix B of this thesis [6].

The TCS will also permit the creation of an MLS network interface by enabling the MYSEA server to enforce all CSLN security decisions outside of the trusted domain of the XTS-400 STOP. The TCS will enforce these decisions differently based upon

whether the communication it receives is inbound to the MYSEA server or outbound to the TCM.

2.3.1.1 TCS Inbound Processing

For inbound connections, the PCC protocol handler will extract key data fields from the deconstruction of each IPsec packet and send that data to the TCS. These fields consist of the source IP addresses from the inner and outer IPsec packet header. The source IP address of the outer header is that of the TCM, which through the *trusted channel*, conveys the packets implicit *sensitivity level*. The TCS will use the IP address of the TCM to query the TCDB and use the returned data to associate an explicit *sensitivity level* to the incoming connection. The TCS will create a new record in the SCIDB and store the explicit *sensitivity level* of each connection and source IP address from the inner header for later use by the SCS. Once the values have been stored in the SCIDB, the TCS shall signal the PCC protocol handler to continue processing the incoming packet.

An example of an AH/ESP *security bundled* packet displaying its primary fields is depicted in Figure 1. The new IP header field is created by the tunnel mode of IPsec and contains the IP address of the tunnel endpoints. For MYSEA, the endpoints are the MYSEA server and its associated TCM. The AH field is specific to the AH protocol and defines the security parameters necessary for each end of the tunnel to process the packet. Likewise, the ESP field defines the security parameters necessary for each end of the tunnel to process the ESP portion of the packet. The original IP header field contains the source and destination IP address from the original IP packet. The next field denotes the layer four protocol of the original packet – TCP for this example. The data field is the data from the original packet.




Figure 1. IPsec AH/ESP Packet

2.3.1.1.1 Initial TCS processing

The PCC protocol handler will first process the AH and ESP protected packet and verify that the packet meets the requirements of the SPD and that the packet correctly decrypts using the prescribed cryptographic algorithms. Upon successful completion of the cryptographic checks, the PCC protocol handler will pass the new and original source IP addresses to the TCS. The TCS will use the source IP address (the IP address of the TCM) from the outer header to query the TCDB and derive the explicit *sensitivity level* of the packet.

2.3.1.1.2 Final TCS Processing

The TCS will create a new record in the SCIDB and write both the original source IP address and the explicit *sensitivity level* to that record for use by the SCS. The TCS will then signal the PCC protocol handler to continue processing the packet. The PCC protocol handler will forward the original IP packet up the networking stack for further processing and eventual action by the SCS. The original IP packet is depicted in Figure 2.




Figure 2. Original IP Packet

2.3.1.2 TCS Outbound Processing

For outbound connection requests, the SCS requests the TCS to verify that the *sensitivity level* of the outbound connection to the requested CSLN is equal to the *sensitivity level* of the requesting application. Upon receipt of the SCS request, the TCS will use the requested destination IP address of the connection to query the SCODB for the explicit *sensitivity level* of the requesting application. Once the explicit *sensitivity level* is obtained, the TCS will use the destination IP address to query the TCDB for the explicit *sensitivity level* of the destination single-level network. The TCS will verify that the sensitivity level of the outgoing connection is equal to the sensitivity level of the destination single-level network. Once the connection is validated, the TCS will signal the SCS to continue processing the connection request through the networking stack.

2.3.1.3 Connection Termination

The TCS shall be responsible for connection teardown. Upon recognition that a connection is terminating between an untrusted MYSEA server process and the CSLN or upon timeout of the connection, the TCS shall delete the appropriate record associated with the connection from its corresponding database, SCIDB or SCODB, and send the PCC protocol handler a delete command to terminate its connection between the MYSEA server and the corresponding TCM.

2.3.2 Secure Connection Server

The SCS is a trusted daemon on the MYSEA server that is placed between the untrusted application protocol servers and the IP layer. For a connection initiated by a single-level network, the SCS process will be started by the TCS upon the receipt of a new connection request. For a connection started by a MYSEA LAN client, the SCS will be started by an existing MYSEA server daemon known as the “Trusted Path Server” [1]. This daemon is responsible for initializing secure operations for each user logged into the MYSEA LAN. The SCS provides two primary functions: one for inbound and the other for outbound connections.

2.3.2.1 Inbound Processing

The SCS maintains an open socket listening for incoming connections already processed by the TCS. Once a connection is accepted, the SCS spawns a child process responsible for handling that connection. The child process shall use the source IP address of the packet to query the SCIDB for the explicit *sensitivity level* of the connection. The child process shall then spawn the requesting application protocol server at the explicit *sensitivity level* returned by the SCIDB query

2.3.2.2 Outbound Processing

The SCS also processes outbound traffic from the Application Protocol servers to the CSLNs. Once a connection request is received, the SCS spawns a trusted child process responsible for that connection. Similar to the PCC protocol handler, the child SCS will extract key data fields from the connection request and post that data to the SCODB for use by the TCS. The child SCS shall create a new record in the SCODB for each new connection request. It will be responsible for writing the requested destination IP address and the *sensitivity level* of the requesting process to the record.

Once complete, the child SCS will request the TCS to verify that the *sensitivity level* of the outbound request is equal to the *sensitivity level* of the destination single-level network. Detailed data flow is provided below.

2.3.2.2.1

The child SCS process will receive the requested connection to a single-level network and create a new record in the SCODB. The child SCS writes the *sensitivity level* of the requesting process and its requested destination IP address to the record.

2.3.2.2.2

If the write is successful, the child SCS process will make a request to the TCS for validation that the requested connection is to an authorized CSLN at the corresponding *sensitivity level*.

2.3.2.2.3

The TCS will use the destination IP address sent from the SCS to query the SCODB for the *sensitivity level* of the application requesting access to the CSLN. The TCS will use the destination IP address to query the TCDB for the *sensitivity level* of the TCM providing security services for the requested CSLN.

2.3.2.2.4

The TCS will then validate that the *sensitivity level* of the requesting application and the *sensitivity level* of the requested CSLN are equal. If equal, the TCS will return authorization to the SCS to continue processing the connection through the networking stack.

2.3.3 Trusted Channel Database

The TCDB is a static database which maintains an association among a TCM, its IP address, and the permitted IP address space for which the TCM provides security services. The integrity of this database is crucial to the overall TSF of the MYSEA server. Thus, all writes to this database shall be constrained only to authorized security administrators. The TCDB shall be administered in accordance with the Δ *Property* [7] as advocated by Fellows. The TCDB shall allow *read-only* access and restrict that access to only the TCS and SCS.

2.3.4 Secure Connection Inbound Database

The SCIDB is a dynamic database which is read and modified by the TCS and is read by the SCS. No other processes are permitted to use the SCIDB interface. Once again, the integrity of this database is crucial to the overall TSF of the MYSEA server. The SCIDB tracks information required to associate the explicit *sensitivity level* of a CSLN to an untrusted application server process. For inbound connections, each record includes the original source IP address and the explicit *sensitivity level* associated with the TCM. Upon connection teardown, the record in the SCIDB shall be deleted by the TCS.

2.3.5 Secure Connection Outbound Database

The SCODB is a dynamic database which is read and modified by both the SCS and TCS. Again, the integrity of this database is crucial to the overall TSF of the MYSEA server. The SCODB tracks information required to associate the *sensitivity level* of an untrusted application to a CSLN. For outbound connections, these fields consist of the requested destination IP address and the *sensitivity level* of the requesting untrusted process. Upon connection teardown, the record in the SCODB shall be deleted by the TCS.

2.4 TRUSTED CHANNEL MODULE

2.4.1 Trusted Channel Server

The TCM TCS mirrors the functionality of the TCS on the MYSEA server for initiating, managing and terminating each PCC. The TCS will also utilize the MYSEA Security Association protocol to bind all *trusted channel* connections between the TCM and MYSEA server.

2.4.2 Network Address Translation Server

The TCM will also provide a Network Address Translation (NAT) server. The requirement for NAT was previously discussed in Section 2.1. The TCM will perform dynamic NAT on all inbound packets and static NAT on all outbound packets.

3. TRUSTED CHANNEL MODULE TO MYSEA SERVER MANAGEMENT FUNCTIONS

3.1 OVERVIEW

The TCM provides the only path between the MYSEA server and one of its multiplexed CSLNs. The TCM permits the MYSEA server to assign an explicit *sensitivity level* to all connections coming from a CSLN by attaching an implied *sensitivity level* to each incoming connection. The PCC enables this functionality by positively authenticating each TCM and by extension, the CSLN for which it provides services. By positively authenticating a TCM and its corresponding CSLN, an implied *sensitivity level* can be associated with the Internet Protocol (IP) routing address of the CSLN interface of the TCM.

The NAT functionality incorporated into the TCM will provide the interface between the segmented IP space of each domain and the IP space used by the multilevel network. The TCM shall serve as each single-level networks interface to the MYSEA servers. As a result, the IP address of the TCM will be advertised in the single-level network in place of the IP address of the MYSEA server. The TCM shall perform dynamic NAT [8] [9] for all inbound packets from the CSLN to the MYSEA server by translating the destination address of each packet from the TCM to the requested MYSEA server.

A data flow analysis is presented in Figure 5.

3.2 TCM INBOUND STATES

The TCM will change states based upon the incoming connections it receives, its authorization to perform the requested NAT function and the security policy decisions made by the Security Policy Database (SPD) of the PCC protocol handler. The following sections will discuss the various states of the TCM.

3.2.1 TCM STATE VARIABLES

The TCM includes three separate variables as shown in Table 1. “Power” indicates that the TCM is either un-powered and dormant or powered and active. “NAT” indicates successful translation from the restricted domain IP address block of the CSLN to the IP address block of the MYSEA server. “Trusted Channel Operation” indicates an

established PCC with the MYSEA server. These three variables provide 2^3 or eight possible states, of which, only four states are reachable.

| Description | Values | Abbreviation |
|-----------------------------|---------------|---------------------|
| Power | On/Off | Power |
| Network Address Translation | Yes/No | NAT |
| Trusted Channel Operation | Yes/No | TCO |

Table 1. TCM State Variables

3.2.2 TCM DISALLOWED STATES

Four states are disallowed by the TCM as shown in Table 2. In other words, no possibility exists to transition into these states.

| Power | NAT | TCO | Reason for Disallowed State |
|--------------|------------|------------|------------------------------------|
| Off | Yes | No | No Power |
| Off | No | Yes | No Power |
| Off | Yes | Yes | No Power |
| On | No | Yes | No TCO operations without NAT |

Table 2. TCM Disallowed States

3.2.3 TCM ALLOWABLE STATES

Four states are allowed by the TCM as show by Table 3 and illustrated in Figure 3.

| State Number | Power | NAT | TCO | Name |
|--------------|-------|-----|-----|-----------|
| 0 | Off | No | No | Power Off |
| 1 | On | No | No | Idle |
| 2 | On | Yes | No | NAT |
| 3 | On | Yes | Yes | TCO |

Table 3. TCM Allowable States

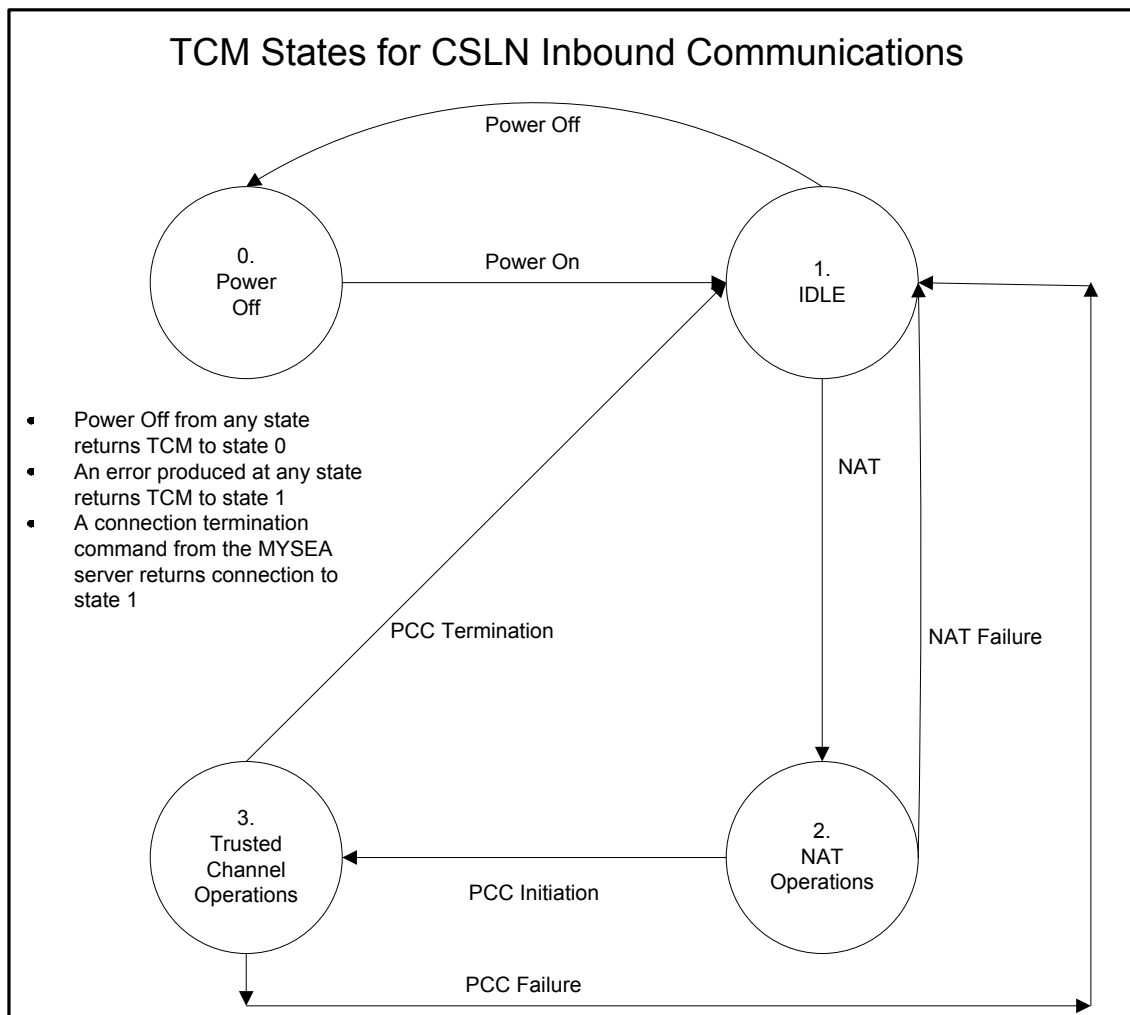


Figure 3. TCM Allowable Inbound State Diagram

3.3 MYSEA SERVER CONNECTED SINGLE LEVEL NETWORK INBOUND STATES

The MYSEA server has a large number of associated states when its complete TSF is considered. This section will be restricted to both the disallowed and allowed states associated with its CSLN operations.

3.3.1 MYSEA SERVER CONNECTED SINGLE LEVEL NETWORK STATES VARIABLES

The MYSEA server CSLN functionality includes four separate variables as shown in Table 4. “Power” indicates that the MYSEA server is either un-powered and dormant or powered and active. “PCC Connected” represents initial IPsec connections between the MYSEA server and TCM. “PCC Authenticated” represents a completed PCC and the successful association of the implicit *sensitivity level* of the connection to its explicit *sensitivity level*. A successful association of the explicit *sensitivity level* consists of a successful TCS write of the inner PCC IP source address to the SCIDB, a trusted read of the *sensitivity level* of the TCM from the TCDB, followed by a trusted write of the *sensitivity level* of the TCM to the SCIDB. “Trusted Operations” represents the SCS trusted read of the *sensitivity level* of the incoming communication from the SCIDB and the successful spawning of the requested application at the correct *sensitivity level*. These four variables provide 2⁴ or sixteen possible states, of which, only five states are reachable.

| Description | Values | Abbreviation |
|--------------------|--------|---------------|
| Power | On/Off | Power |
| PCC Connected | Yes/No | Connected |
| PCC Authenticated | Yes/No | Authenticated |
| Trusted Operations | Yes/No | Trusted |

Table 4. MYSEA Server CSLN State Variables

3.3.2 MYSEA SERVER CONNECTED SINGLE LEVEL NETWORK DISALLOWED STATES

The following eleven states are disallowed by the MYSEA server CSLN operations as shown in Table 5. In other words, no possibility exists to transition into these states.

| Power | Connected | Authenticated | Trusted | Reason |
|-------|-----------|---------------|---------|--|
| Off | N | N | Y | No Power |
| Off | N | Y | N | No Power |
| Off | Y | N | N | No Power |
| Off | N | Y | Y | No Power |
| Off | Y | N | Y | No Power |
| Off | Y | Y | N | No Power |
| Off | Y | Y | Y | No Power |
| On | N | N | Y | No Trusted without Connected and Authenticated |
| On | N | Y | N | No Authenticated without Connected |
| On | N | Y | Y | No Authenticated and Trusted without Connected |
| On | Y | N | Y | No Trusted without Authenticated |

Table 5. MYSEA Server CSLN Disallowed States

3.3.3 MYSEA SERVER CONNECTED SINGLE LEVEL NETWORK ALLOWED STATES

The following five states are allowed by the MYSEA server CSLN as show by Table 6 and illustrated in Figure 5.

| State Number | Power | Connected | Authenticated | Trusted | Reason |
|--------------|-------|-----------|---------------|---------|---------------|
| 0 | Off | N | N | N | Power Off |
| 1 | On | N | N | N | Idle |
| 2 | On | Y | N | N | Connected |
| 3 | On | Y | Y | N | Authenticated |
| 4 | On | Y | Y | Y | Trusted |

Table 6. MYSEA Server CSLN Allowed States

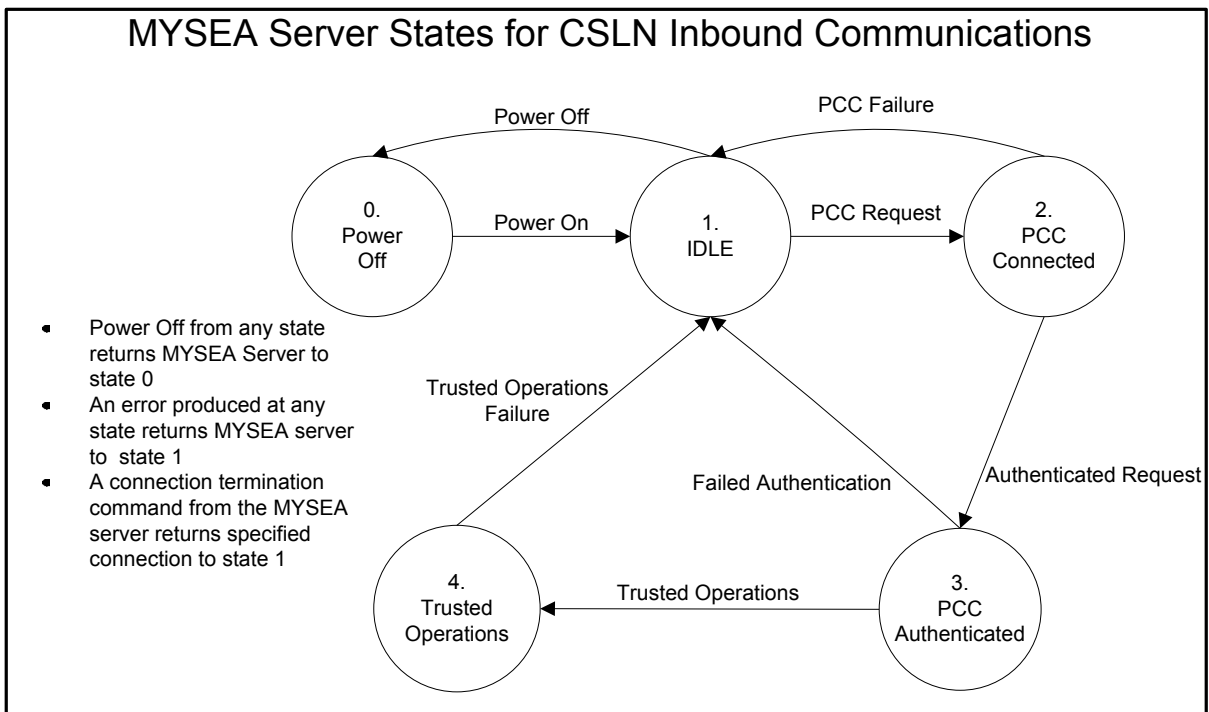


Figure 4. MYSEA Server Inbound Allowable State Diagram

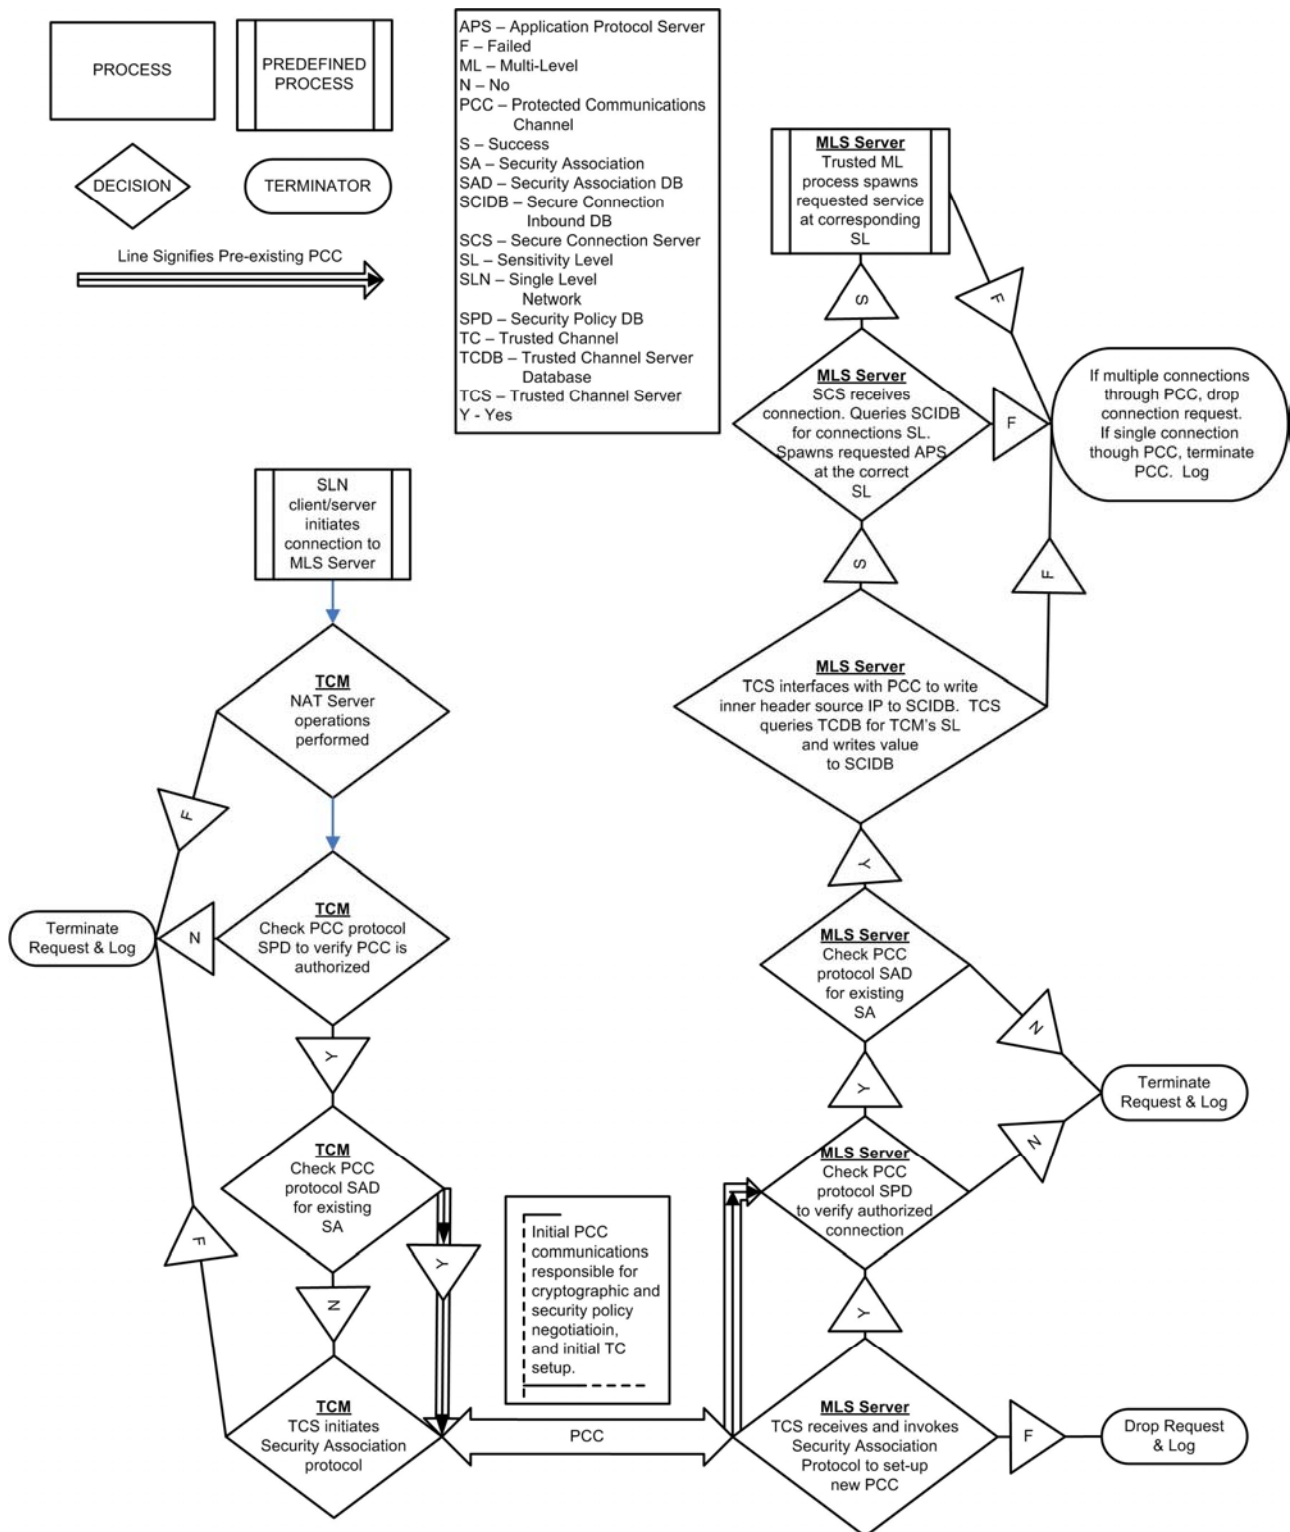


Figure 5. TCM to MYSEA Server Flow Chart

THIS PAGE INTENTIONALLY LEFT BLANK

4. MYSEA SERVER TO TRUSTED CHANNEL MODULE MANAGEMENT FUNCTIONS

4.1 OVERVIEW

The MYSEA server and TCM sustain a peer-to-peer relationship for initiating connections. The MYSEA server shall retain the capability to initiate connections through dedicated single level network interfaces. Additionally, the MYSEA server shall provide the capability to initiate connections through MLS network interfaces. These interfaces shall be capable of multiplexing a large number of connections at multiple *sensitivity levels* to a large number of TCMs. The creation of a *trusted channel* is the key component required for the MYSEA server to provide an MLS network interface. The MYSEA server invokes a mandatory access control (MAC) check to ensure the *sensitivity level* of the process requesting the outbound connection is equivalent to the *sensitivity level* of the intended CSLN. Once the MAC check is completed, the PCC is initiated to provide the required *trusted channel* necessary to provide a connection through an MLS network interface. Due to the aforementioned DoD IP space restrictions, the TCM will perform static source NAT [8] [9] for all outbound connections from the MYSEA server to the CSLN.

The TCM shall perform static NAT for all outbound communication from the MYSEA server to the CSLN by conducting a one to one swap of the source IP address from the MYSEA server to that of the source IP address of the TCM. A data flow analysis is presented in Figure 8.

4.2 MYSEA SERVER CONNECTED SINGLE LEVEL NETWORK OUTBOUND STATES

The CSLN operations of the MYSEA server will change states based upon its outbound connections and the decisions made with respect to those connections by its PCC SPD. The following sections will discuss the various states of the MYSEA server.

4.2.1 MYSEA SERVER CONNECTED SINGLE LEVEL NETWORK STATES VARIABLES

The MYSEA server CSLN functionality includes four separate variables as shown in Table 7. “Power” indicates that the MYSEA server is either un-powered and dormant or powered and active. “Trusted CSLN Operations Authorized” represents an untrusted

process accessing the SCS to initiate a valid CSLN connection. The SCS shall write the *sensitivity level* of the process and requested destination IP address to the SCODB. The SCS shall then request the TCS to query the SCODB with the destination IP address to obtain the explicit *sensitivity level* of the requesting application protocol server. The TCS shall then query the TCDB with the destination IP address of the requested connection to obtain the explicit *sensitivity level* of the destination TCM. The TCS shall perform a MAC equivalence check and verify that the *sensitivity level* of the server is equal to the *sensitivity level* of the requested CSLN. “PCC Connected” represents initial IPsec communications between the MYSEA server and TCM. “*Trusted Channel* Authenticated” represents an authenticated PCC. These four variables provide 2^4 or sixteen possible states, of which, only five states are reachable.

4.2.2 MYSEA SERVER CONNECTED SINGLE LEVEL NETWORK DISALLOWED STATES

Eleven states are disallowed by the MYSEA server CSLN operations as shown in Table 8. In other words, no possibility exists to transition into these states.

| Description | Values | Abbreviation |
|------------------------------------|--------|---------------|
| Power | On/Off | Power |
| Trusted CSLN Operations Authorized | Yes/No | Authorized |
| PCC Connected | Yes/No | Connected |
| Trusted Channel Authenticated | Yes/No | Authenticated |

Table 7. MYSEA Server CSLN State Variables

| Power | Authorized | Connected | Authenticated | Reason |
|-------|------------|-----------|---------------|---|
| Off | N | N | Y | No Power |
| Off | N | Y | N | No Power |
| Off | Y | N | N | No Power |
| Off | N | Y | Y | No Power |
| Off | Y | N | Y | No Power |
| Off | Y | Y | N | No Power |
| Off | Y | Y | Y | No Power |
| On | N | N | Y | No Authenticated without Connected and Authorized |
| On | N | Y | N | No Connected without Authorized |
| On | N | Y | Y | No Connected and Authenticated without Authorized |
| On | Y | N | Y | No Authenticated without Connected |

Table 8. MYSEA Server CSLN Disallowed States

4.2.3 MYSEA SERVER CONNECTED SINGLE LEVEL NETWORK ALLOWED STATES

The following five states are allowed by the MYSEA server CSLN as show by Table 9 and illustrated in Figure 6.

| State Number | Power | Authorized | Connected | Authenticated | Reason |
|--------------|-------|------------|-----------|---------------|---------------|
| 0 | Off | N | N | N | Power Off |
| 1 | On | N | N | N | Idle |
| 2 | On | Y | N | N | Authorized |
| 3 | On | Y | Y | N | Connected |
| 4 | On | Y | Y | Y | Authenticated |

Table 9. MYSEA Server CSLN Allowed States

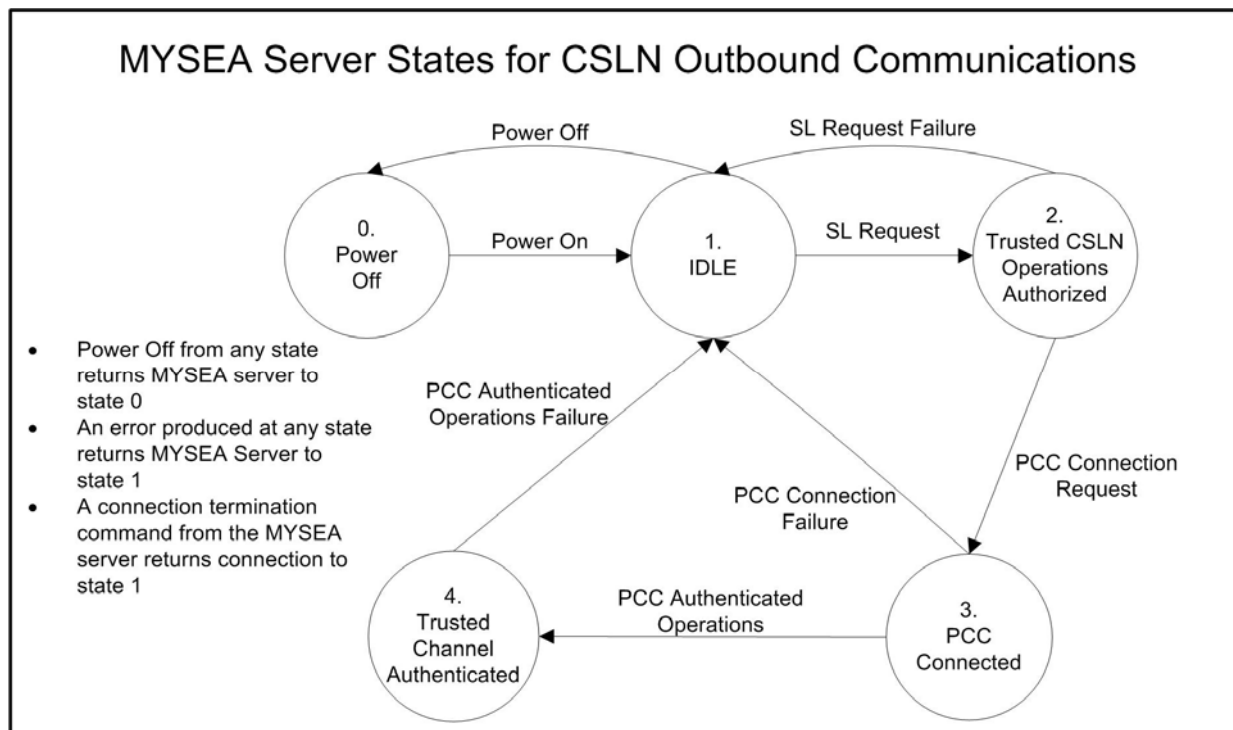


Figure 6. MYSEA Server Outbound Allowable States

4.3 TCM OUTBOUND STATES

The TCM will change states based on its ability to negotiate and operate a successful PCC and on its ability to perform NAT operations on the incoming packet. The following sections will discuss the various states of the TCM.

4.3.1 TCM STATE VARIABLES

The TCM includes three separate variables as shown in Table 10. “Power” indicates that the TCM is either un-powered and dormant or powered and active. “Trusted Channel Operation” indicates an established PCC with the MYSEA server. “NAT” indicates a successful source NAT from the IP address of the MYSEA server to the IP address of the TCM. These three variables provide 2^3 or eight possible states, of which, only four states are reachable.

| Description | Values | Abbreviation |
|-----------------------------|--------|--------------|
| Power | On/Off | Power |
| Trusted Channel Operation | Yes/No | TCO |
| Network Address Translation | Yes/No | NAT |

Table 10. TCM State Variables

4.3.2 TCM DISALLOWED STATES

The following four states are disallowed by the TCM as shown in Table 11. In other words, no possibility exists to transition into these states.

| Power | TCO | NAT | Reason for Disallowed State |
|-------|-----|-----|-------------------------------|
| Off | Yes | No | No Power |
| Off | No | Yes | No Power |
| Off | Yes | Yes | No Power |
| On | No | Yes | No NAT operations without TCO |

Table 11. TCM Disallowed States

4.3.3 TCM ALLOWABLE STATES

The following four states are allowed by the TCM as show by Table 12 and illustrated in Figure 7.

| State Number | Power | TCO | NAT | Name |
|--------------|-------|-----|-----|--------------------|
| 0 | Off | No | No | Power Off |
| 1 | On | No | No | Idle |
| 2 | On | Yes | No | Trusted Operations |
| 3 | On | Yes | Yes | NAT Operations |

Table 12. TCM Outbound Allowable States

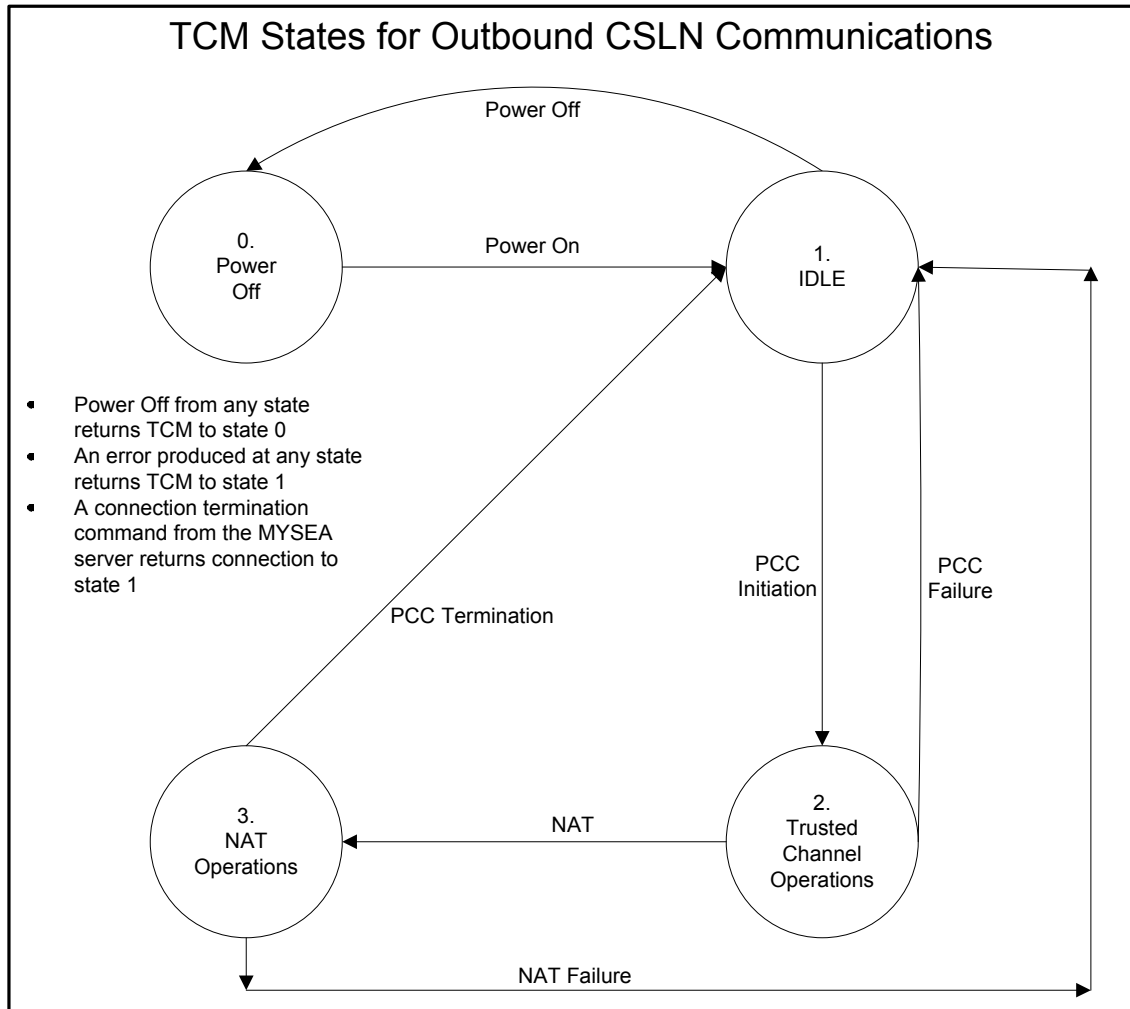


Figure 7. TCM Outbound Allowable States

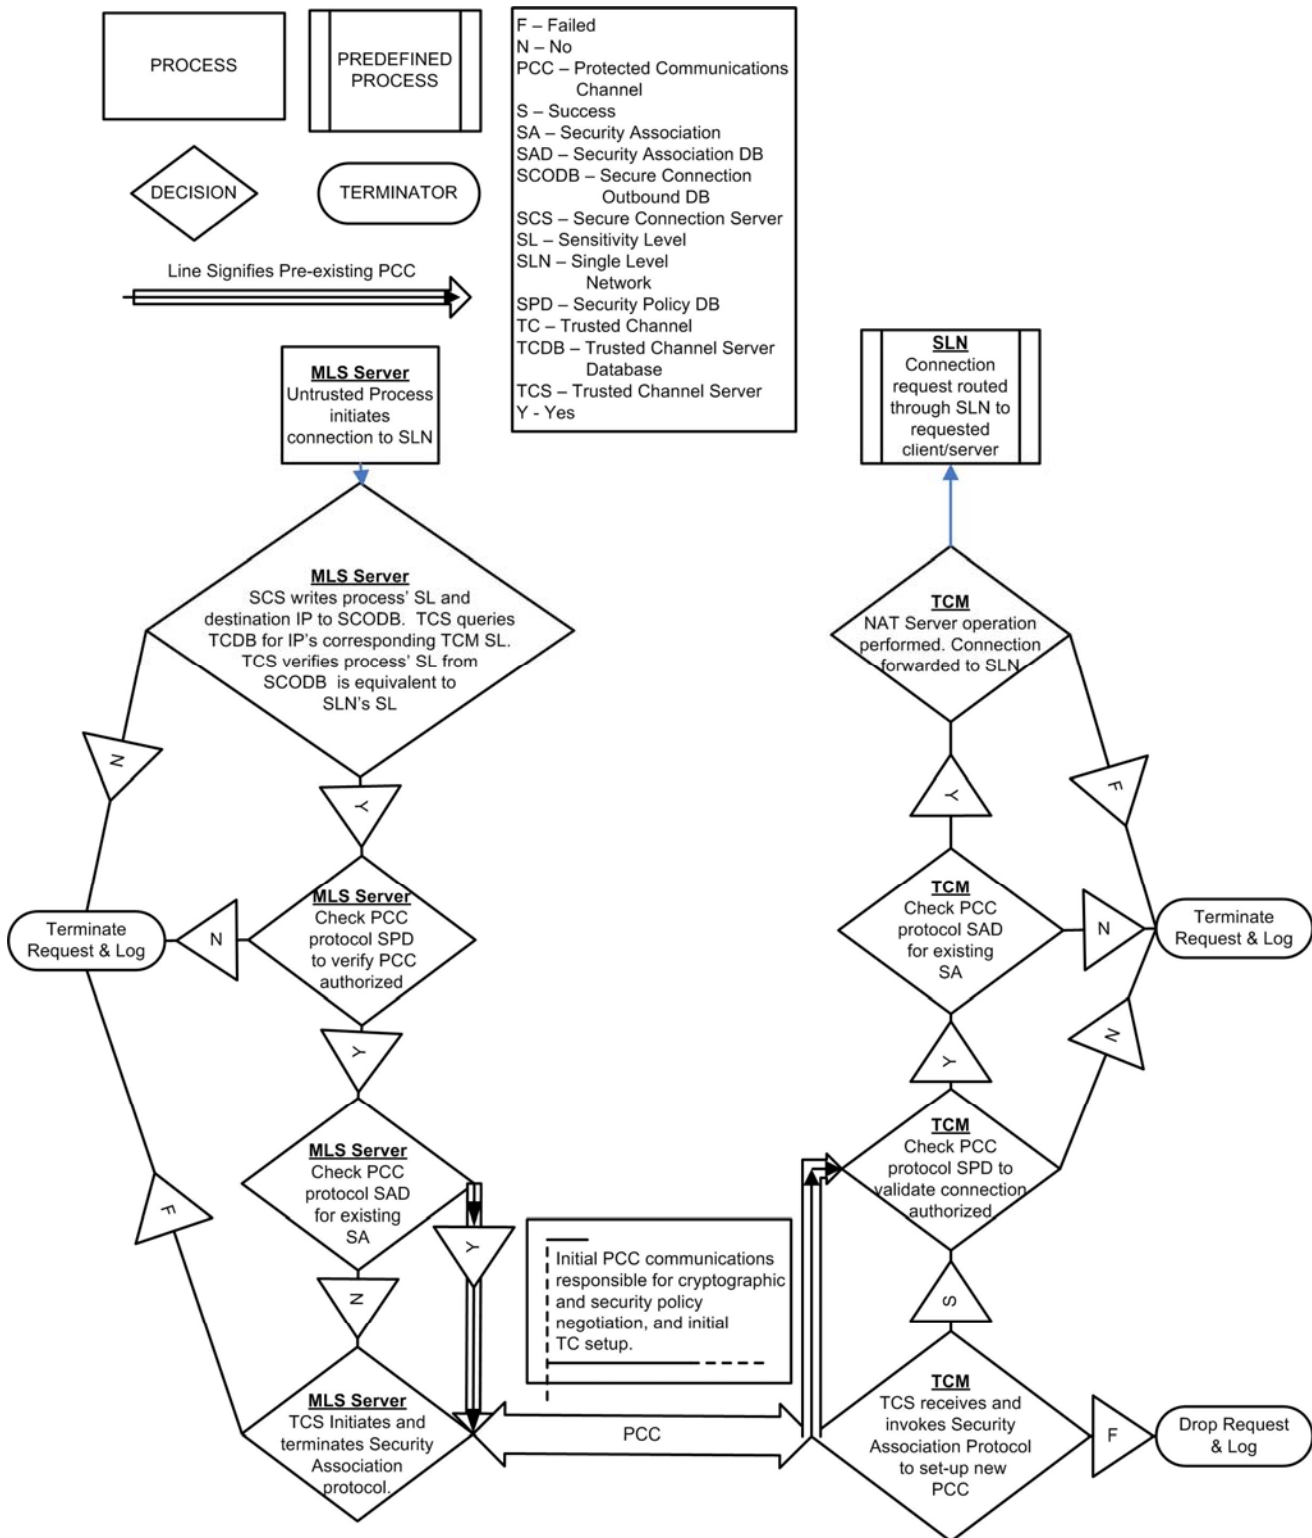


Figure 8. MYSEA Server to TCM Flow Chart

APPENDIXES

APPENDIX A ABBREVIATIONS, ACRONYMS, AND DEFINITIONS

A.1 Abbreviations, Acronyms

| | |
|---------|--|
| AH | IP Authentication Header Protocol |
| CSLN | MYSEA Connected Single Level Network(s) |
| DoD | Department of Defense |
| ESP | IP Encapsulating Security Payload Protocol |
| JWICS | Joint Worldwide Intelligence Communications System |
| MAC | Mandatory Access Control |
| MLS | Multilevel Security |
| MYSEA | Monterey Security Architecture |
| NAT | Network Address Translation |
| NIPRNET | NonSecure Internet Protocol Router Network |
| NIC | Network Information Center |
| PCC | Protected Communications Channel |
| SCIDB | Secure Connection Inbound Database |
| SCODB | Secure Connection Outbound Database |
| SIPRNET | SECRET Internet Protocol Router Network |
| SCS | Secure Connection Server |
| TCM | Trusted Channel Module |
| TCS | Trusted Channel Server |
| TCDB | Trusted Channel Database |
| TSF | Target of Evaluation Security Function |
| TSP | Target of Evaluation Security Policy |

A.2 Definitions

2.12 Δ Property: Requires all security-related modification to the MLS distributed architecture to be made only by explicitly authorized trusted agents [7].

2.13 Connected Single Level Network (CSLN): The segment of the MYSEA architecture from the MYSEA server to its supported TCMs responsible for providing single level network connectivity to the MYSEA server.

2.14 Security Bundle: An IPsec security association that invokes both the Authentication Header protocol and the IP Encapsulating Security Payload protocol onto one IP packet [10].

2.15 Sensitivity Level: The combined classification of data based upon its security or classification level and integrity level.

2.16 Target of Evaluation (TOE): An IT product or system and its associated guidance documentation that is the subject of an evaluation [3].

2.17 TOE Security Functions (TSF): A set consisting of all hardware, software, and firmware of the TOE that must be relied upon for the correct enforcement of the TSP [3].

2.18 TOE Security Policy (TSP): A set of rules that regulate how assets are managed, protected and distributed within a TOE [3].

2.19 Trusted channel: A means by which a TSF and a remote trusted IT product can communicate with necessary confidence to support the TSP [3].

2.20 Trusted Channel Module: Security device required to enable “Inter-TSF Trusted Channels” between the MYSEA server and its authorized single level network.

2.21 Tunnel Mode: A security association applied to an IP tunnel, with the access controls applied to the headers of the traffic inside the tunnel [10].

APPENDIX B REFERENCES

- [1] J. D. Wilson, "A Trusted Connection Framework for Multilevel Secure Local Area Networks", Naval Postgraduate School, June 2000.
- [2] C. E. Irvine, T. E. Levin, T. D. Nguyen, D. Shifflett, J. Khosalim, P. C. Clark, A. Wong, F. Afinidad, D. Bibighaus, and J. Sears, "Overview of a High Assurance Architecture for Distributed Multilevel Security", Proceedings of the 2004 IEEE Systems Man and Cybernetics Information Assurance Workshop, West Point, NY, June 2004, pp. 38-45.
- [3] "Common Criteria for Information Technology Security Evaluation Volume I", *CCIMB-2004-01-001*, Version 2.2 ed: International Organization for Standardization, January 2004.
- [4] DigitalNet Government Solutions LLC, "XTS-400 Trusted Computer System Technical Overview", Herndon, VA, July 2004,
http://www.digitalnet.com/solutions/info_sec_sol/xts400_trusted_sys.htm,
15 July 04.
- [5] Defense Information Systems Agency, "SIPRNET Classification Guide".
- [6] J. D. Sears, "Protected Communications Channel Protocol Requirements Document," Appendix B, "Simultaneous Connection Management and Protection in a Distributed Multilevel Security Environment", Masters Thesis, Naval Postgraduate School, Monterey, California, September 2004.
- [7] J. Fellows, J. Hemenway, N. Kelem, and S. Romero, "The Architecture of a Distributed Trusted Computing Base", *Proceedings of the 10th National Conference on Computer Security*, pp. 68-77, September 1987.
- [8] Cisco Systems Inc "How NAT Works, Document ID: 6450",
http://www.cisco.com/en/US/tech/tk648/tk361/technologies_tech_note09186a0080094831.shtml, 5 August 2004.
- [9] K. Srisuresh and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT), *RFC 3022*", *Network Working Group, IETF*, January 2001.
- [10] S. Kent and K. Seo, "Security Architecture for the Internet Protocol, Draft-IETF-IPsec-RFC2401BIS-01", *IP Security Protocol Working Group, IETF*, July 2004.

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF REFERENCES

- [1] Department of Defense, "Information Assurance," ASD/C3I, DODD 8500.1, October 2002.
- [2] "Common Criteria for Information Technology Security Evaluation Volume I", *CCIMB-2004-01-001*, Version 2.2 ed: International Organization for Standardization, January 2004.
- [3] "Common Criteria for Information Technology Security Evaluation Volume II", *CCIMB-2004-01-002*, Version 2.2 ed: International Organization for Standardization, January 2004.
- [4] Department of Defense, "Trusted Network Interpretation of the TCSEC", National Computer Security Center, NCSC-TG-005, July 1987.
- [5] J. D. Wilson, "A Trusted Connection Framework for Multilevel Secure Local Area Networks", Naval Postgraduate School, June 2000.
- [6] C. E. Irvine, D. Shifflett, P. C. Clark, T. E. Levin, and G. W. Dinolt, "MYSEA Security Architecture," Naval Postgraduate School Center for Information System Security Studies and Research, NPS-CS-02-006, May 2002.
- [7] C. E. Irvine, T. E. Levin, J. D. Wilson, D. Shifflett, and B. Pereira, "An Approach to Security Requirements Engineering for a High Assurance System", *Requirements Engineering*, Vol. 7. No. 4, May 2002, pp. 192-208.
- [8] C. E. Irvine, T. E. Levin, T. D. Nguyen, D. Shifflett, J. Khosalim, P. C. Clark, A. Wong, F. Afinidad, D. Bibighaus, and J. Sears, "Overview of a High Assurance Architecture for Distributed Multilevel Security", Proceedings of the 2004 IEEE Systems Man and Cybernetics Information Assurance Workshop, West Point, NY, June 2004, pp. 38-45.
- [9] DigitalNet Government Solutions LLC, "XTS-400 Trusted Computer System Technical Overview", Herndon, VA, July 2004,
http://www.digitalnet.com/solutions/info_sec_sol/xts400_trusted_sys.htm,
15 July 04.
- [10] Department of Defense, "Trusted Computer System Evaluation Criteria", National Computer Security Center, DoD 5200.28-STD, December 1985.
- [11] D. E. Bell and L. LaPadula, "Secure Computer System: Unified Exposition and Multics Interpretation", Mitre Corp., Hanscomb AFB, MA, Tech Rep ESD-TR-76-372, 1975.

- [12] K. J. Biba, "Integrity Considerations for Secure Computer Systems", MITRE Corp., Tech. Rep. ESD-TR-76-372, 1977.
- [13] Irvine, Cynthia E., Levin, Timothy E., Nguyen, Thuy D., and Dinolt, George W., "The Trusted Computing Exemplar Project", Proceedings of the 2004 IEEE Systems Man and Cybernetics Information Assurance Workshop, West Point, NY, June 2004, pp. 109 - 115.
- [14] J. C. Williams and M. L. Day, "Sensitivity Labels and Security Profiles", Proceedings of the Eleventh National Computer Security Profiles, October 1988.
- [15] S. B. Lipner, "Non-Discretionary Controls for Commercial Applications", *1982 IEEE Symposium on Security and Privacy, Oakland, CA*, April 1982.
- [16] J. Fellows, J. Hemenway, N. Kelem, and S. Romero, "The Architecture of a Distributed Trusted Computing Base", *Proceedings of the 10th National Conference on Computer Security*, pp. 68-77, September 1987.
- [17] C. Weissman, "BLACKER: Security for the DDN, Examples of A1 Security Engineering Trades", 1992 IEEE Computer Society symposium on Research in Security and Privacy, Oakland, CA, May 1992, pp. 286-292.
- [18] M. Vetterling, G. Wimmel, and A. Wisspeintner, "Secure Systems Development Based on the Common Criteria: The PalME Project", ACM SIGSOFT 2002 Software Engineering Notes, Charleston, SC, 2002, pp. 129-138.
- [19] National Security Agency Information Assurance Directorate, "Consistency Instruction Manual for Development of US Government Protection Profiles (PP) for use in Medium Robustness Environments", Release 2.0, March 2004.
- [20] Cisco Systems Inc., "Security Target for Cisco IOS/IPSEC", Version: 3.7, San Jose, CA, 16 September 2002.
- [21] National Security Agency Information Assurance Directorate, "U.S. Government Protection Profile for Separation Kernels in Environments Requiring High Robustness", Version 0.621, July 2004.
- [22] J. D. Sears, "Connected Single Level Network System Requirements Document", Appendix A, "Simultaneous Connection Management and Protection in a Distributed Multilevel Security Environment", Masters Thesis, Naval Postgraduate School, Monterey, California September 2004.
- [23] J. D. Sears, "Connected Single Level Network Trusted Channel Management Requirements Document", Appendix C, "Simultaneous Connection Management and Protection in a Distributed Multilevel Security Environment", Masters Thesis, Naval Postgraduate School, Monterey, California, September 2004.

- [24] K. Srisuresh and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT), RFC 3022", *Network Working Group, IETF*, January 2001.
- [25] C. Kaufman, "The Internet Key Exchange (IKEv2) Protocol, draft-ietf-ipsec-ikev2-11.txt," *IP Security Protocol Working Group, IETF*, October 2003.
- [26] J. D. Sears, "Protected Communications Channel Protocol Requirements Document", Appendix B, "Simultaneous Connection Management and Protection in a Distributed Multilevel Security Environment", Masters Thesis, Naval Postgraduate School, Monterey, California, September 2004.
- [27] S. Kent and K. Seo, "Security Architecture for the Internet Protocol, Draft-IETF-IPsec-RFC2401BIS-01", *IP Security Protocol Working Group, IETF*, July 2004.
- [28] Defense Information Systems Agency, "SIPRNET Classification Guide."
- [29] S. Kent, "IP Encapsulating Security Payload (ESP), RFC 2406", *Network Working Group, IETF*, November 1998.
- [30] S. Kent, "IP Authentication Header, RFC 2402", *Network Working Group, IETF*, November 1998.
- [31] National Institute of Standards and Technology, "Federal Information Processing Standards Publication 197 Announcing the Advanced Encryption Standard", <http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf>, November 2001.
- [32] National Institute of Standards and Technology, "Federal Information Processing Standards Publication 180-2 Announcing the Secure Hash Standard", <http://csrc.nist.gov/publications/fips/fips180-2/fips180-2.pdf>, August 2002.
- [33] Department of Defense, "Directive C-5200.5 Communications Security", 21 April 1990.
- [34] T. Li, B. Cole, P. Morton, and D. Li, "Cisco Hot Standby Router Protocol (HSRP), RFC 2281", *Network Working Group, IETF*, March 1998.

THIS PAGE INTENTIONALLY LEFT BLANK

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
Ft. Belvoir, VA
2. Dudley Knox Library
Naval Postgraduate School
Monterey, CA
3. Susan Alexander
National Security Agency
Fort Meade, MD
4. George Bieber
Office of Secretary of Defense
Washington, DC
5. RADM Joseph Burns
Commander Naval Security Group
Fort George Meade, MD
6. CAPT Mike Brown
Naval Information Warfare Activity
FT George Meade, MD
7. Deborah Cooper
DC Associates, LLC
Roslyn, VA
8. CAPT Timothy V. Flynn
SPAWAR Systems Center, San Diego
San Diego, CA
9. CDR Daniel L. Currie
SPAWAR PMW 161
San Diego, CA
10. Rita Painter
SPAWAR, Code 2721
Monterey, CA

11. LCDR James Downey
NAVSEA
Washington, DC
12. Mike Focke
DigitalNet
Herndon, VA
13. Dr. Diana Gant
National Science Foundation
14. Jennifer Guild
SPAWAR Systems Center, Charleston
Charleston, SC
15. Richard Hale
Defense Information Systems Agency
Falls Church, VA
16. LCDR Scott D. Heller
SPAWAR, PMW-161
San Diego, CA
17. Wiley Jones
Office of Secretary of Defense
Washington, DC
18. Russell Jones
OPNAV N641
Arlington, VA
19. David Ladd
Microsoft Corporation
Redmond, WA
20. Dr. Carl Landwehr
National Science Foundation
Arlington, VA
21. Steve LaFountain
National Security Agency
Fort Meade, MD
22. Dr. Greg Larson
IDA
Alexandria, VA

23. Penny Lehtola
National Security Agency
Fort Meade, MD
24. Ernest Lucier
Federal Aviation Administration
Washington, DC
25. CAPT Sheila McCoy
Headquarters U.S. Navy
Arlington, VA
26. Dr. Vic Maconachy
National Security Agency
Fort Meade, MD
27. Doug Maughan
Department of Homeland Security
Washington, DC
28. Dr. John Monastra
Aerospace Corporation
Chantilly, VA
29. John Mildner
SPAWAR Systems Center, Charleston
Charleston, SC
30. Steve Rose
DigitalNet
Herndon, VA
31. P.J. Jenket
SPAWAR Systems Center, Charleston
Charleston, SC
32. Keith Schwalm
Good Harbor Consulting, LLC
Washington, DC
33. Dr. Ralph Wachter
Office of Naval Research
Arlington, VA

34. George Webber
DigitalNet
Herndon, VA
35. David Wirth
OPNAV N641
Arlington, VA
36. Daniel Wolf
National Security Agency
Fort Meade, MD
37. CAPT Robert Zellmann
OPNAV N614
Arlington, VA
38. Mr. Lew Gutman
SPAWAR Systems Center, San Diego
San Diego, CA
39. Dr. Cynthia E. Irvine
Naval Postgraduate School
Monterey, CA
40. Thuy D. Nguyen
Naval Postgraduate School
Monterey, CA
41. LT Joseph D. Sears
Naval Postgraduate School
Monterey, CA